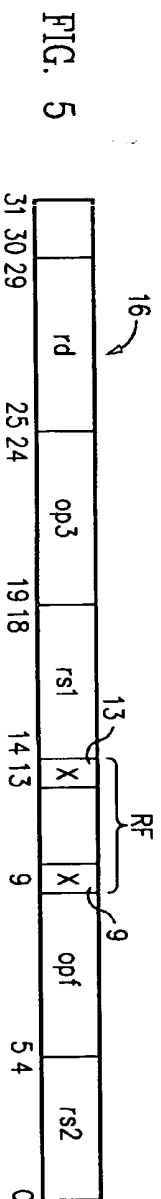
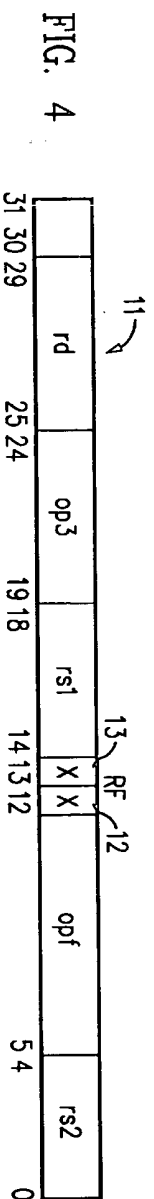


	L #	Hits	Search Text	DBs
1	L4	493	(select\$3 execut\$3) near20 (thread context) near20 (sequential\$3 robin)	USPAT; US-PGPUB
2	L5	1692	(instruction adj2 (pointer counter) ip (address near10 (fetch\$3 prefetch\$3))) near20 (thread context)	USPAT; US-PGPUB
3	L6	75	4 and 5	USPAT; US-PGPUB
4	L9	1	7 and 8	EPO; JPO; DERWENT; IBM_TDB
5	L10	1889	(instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (thread context)	USPAT; US-PGPUB
6	L11	19	4 and 10 not 6	USPAT; US-PGPUB
7	L12	114	(instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (thread context)	EPO; JPO; DERWENT; IBM_TDB
8	L13	2	7 and "13"	EPO; JPO; DERWENT; IBM_TDB
9	L14	3	7 and 12	EPO; JPO; DERWENT; IBM_TDB
10	L8	89	(instruction adj2 (pointer counter) ip (address near10 (fetch\$3 prefetch\$3))) near20 (thread context)	EPO; JPO; DERWENT; IBM_TDB
11	L7	34	(select\$3 execut\$3) near20 (thread context) near20 (sequential\$3 robin)	EPO; JPO; DERWENT; IBM_TDB
12	L20	45	(plural plurality multiple multiplicity several number) adj2 (instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (multithread\$3 thread context) not (6 11)	USPAT; US-PGPUB
13	L21	1	(plural plurality multiple multiplicity several number) adj2 (instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (multithread\$3 thread context) not 14	EPO; JPO; DERWENT; IBM_TDB

opf[3:0]																
opf[8:4]	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
00	—	FMOVb	FMOVd	FMOVq	—	FNEGb	FNEGd	FNEGq	—	FABsb	FABsd	FABsq	—	—	—	—
01	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
02	—	—	—	—	—	—	—	—	—	FSQRts	FSQRtd	FSQRtq	—	—	—	—
03	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
04	—	FADDsb	FADDd	FADDq	—	FSUBsb	FSUBd	FSUBq	—	FMULsb	FMULd	FMULq	—	FDIVsb	FDIVd	FDIVq
05	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
06	—	—	—	—	—	—	—	—	—	FSMULd	—	—	—	—	FSMULq	—
07	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
08	—	FSTOx	FdTOx	FqTOx	FxTOs	—	—	—	FxTOd	—	—	—	FxTOq	—	—	—
09	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
0A	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
0B	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—
0C	—	—	—	—	FIToS	—	FdToS	FqToS	FITod	FsTod	—	FqToD	FIToq	FsToq	FdToq	—
0D	—	FSTOI	FdTOI	FqTOI	—	—	—	—	—	—	—	—	—	—	—	—
0E..1F	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—	—

(PRIOR ART)
FIG. 3B



	Docum ent ID	U	Title	Current OR
1	US 20040 05499 0 A1	<input type="checkbox"/>	Post-pass binary adaptation for software-based speculative precomputation	717/124
2	US 20040 03475 9 A1	<input checked="" type="checkbox"/>	Multi-threaded pipeline with context issue rules	712/1
3	US 20040 03286 5 A1	<input checked="" type="checkbox"/>	Apparatus and method for establishing a call connection state in a packet data communication system	370/367
4	US 20030 21266 0 A1	<input checked="" type="checkbox"/>	Database scattering system	707/1
5	US 20030 18956 5 A1	<input checked="" type="checkbox"/>	Single semiconductor graphics platform system and method with skinning, swizzling and masking capabilities	345/418
6	US 20030 11295 6 A1	<input checked="" type="checkbox"/>	Transferring a call to a backup according to call context	379/221 .01
7	US 20030 11224 6 A1	<input checked="" type="checkbox"/>	Blending system and method in an integrated computer graphics pipeline	345/519
8	US 20030 11224 5 A1	<input checked="" type="checkbox"/>	Single semiconductor graphics platform	345/506
9	US 20030 10305 4 A1	<input checked="" type="checkbox"/>	Integrated graphics processing unit with antialiasing	345/506
10	US 20030 10305 0 A1	<input checked="" type="checkbox"/>	Masking system and method for a graphics processing framework embodied on a single semiconductor platform	345/426
11	US 20030 09754 8 A1	<input checked="" type="checkbox"/>	Context execution in pipelined computer processor	712/228
12	US 20030 06160 1 A1	<input checked="" type="checkbox"/>	Data processing apparatus and method, computer program, information storage medium, parallel operation apparatus, and data processing system	717/144
13	US 20030 04111 0 A1	<input checked="" type="checkbox"/>	System, Method and Structure for generating and using a compressed digital certificate	709/206
14	US 20030 03880 8 A1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for a sequencer in a transform/lighting module capable of processing multiple independent execution threads	345/506
15	US 20030 03722 8 A1	<input checked="" type="checkbox"/>	System and method for instruction level multithreading scheduling in a embedded processor	712/245
16	US 20030 03497 5 A1	<input checked="" type="checkbox"/>	Lighting system and method for a graphics processor	345/426
17	US 20030 02072 0 A1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for a sequencer in a transform/lighting module capable of processing multiple independent execution threads	345/506

invention. Each bundle includes a plurality of instructions and a template field grouped together in a N-bit field. The instructions are located in instruction slots of the N-bit field, with the template field specifying a mapping of the instruction slots to the execution unit types. Other encodings or grouping of instructions are also possible, and are considered within the scope of the present invention.

With continuing reference to the embodiment of FIG. 1, the input to rotator buffer 24 can be zero, one, or two instruction bundles. The number of bundles issued to the EXP stage that are consumed by the execution core may also be zero, one, or two bundles. Note, however, that information about number of bundles available in the number of bundles consumed by the execution core is not available to the rotate pipeline stage logic until the middle of the CPU clock cycle.

Rotate buffer 24 is organized as a circular, wrap-around buffer, which is implemented in the embodiment of FIG. 1 with eight entries, labeled IB₀-IB₇. Rotator buffer 24 has an associated starting write pointer 23 that points to a starting entry for filling (i.e., writing) of instruction bundles into the buffer. Fill operations are controlled by fill control logic block 22. Filling operations take place in-order, that is, according to the original code sequence of the program. After each entry location of buffer 24 gets written with an instruction bundle, fill pointer 23 advances in one position in the buffer stack.

On the readout side, a starting read pointer (i.e., drain pointer) 25 is also associated with buffer 24. FIG. 1 shows the starting read pointer point to IB₀, with up to four instruction bundles (IB₀-IB₃) being read out each clock cycle. The speculative read of the buffer is performed by reading out four consecutive bundles starting from the starting read pointer location, and sending the instruction data to the alignment multiplexers 26 and 27. Depending on the number of bundles consumed by the execution core, the number of bundles available from buffer 24, and the number read from cache 11, the appropriate bundles from the speculative read and the bypass case are selected from each alignment multiplexer.

According to the embodiment of FIG. 1, the number of bundles consumed by the execution core is at most two; the number of bundles read from buffer 24 is at most four; and the number read from cache 11 is at most two. Drain control logic unit 28 receives information about the number of bundles consumed in the execution core. Once the bundle consumption information becomes available, two different bypass combinations and four consecutive bundles from rotator buffer 24 are speculatively read out and then selected. Alignment multiplexers 26 and 27 are an in-order output to the dispersal stages to maintain the FIFO scheme. The alignment multiplexers perform any needed rotation, selecting the appropriate bundles to be latched in dispersal latches 31 and 32, respectively. Each of the alignment multiplexers has six inputs, two of which are the instruction cache way multiplexer (21) outputs, with each bypass instance consisting of two instruction bundles. Also included among the inputs to the alignment multiplexers are the four other consecutive instruction bundles read out from rotator buffer 24.

As the execution engine of the processor consumes instruction bundles, next in-order bundles get consecutively latched into the dispersal latches through the alignment multiplexers. In operation, drain pointer 25 basically trails fill pointer 23, wrapping around the end of the rotate buffer.

FIG. 2 is a block diagram of the execution engine of the processor, showing the various components and their interconnections. The execution engine 200 includes a first logic circuit 201 that divides the cache line into instruction bundles, each of which is written into an entry of the buffer 202, and a second logic circuit 203 that reads out a number of consecutive instruction bundles from the buffer 202, and a buffer 204 having a starting write pointer and a starting read pointer. An execution engine 205 for executing the instructions, an instruction cache 206 that stores a cache line of instructions, and a processor comprising: a second logic circuit 207 that reads out a number of consecutive instruction bundles from the circular buffer 202, and a first logic circuit 208 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 202, and an execution engine 209 for executing the instructions, a circular buffer 210 having a plurality of entries, an instruction cache 211 that stores a cache line of instructions, and a processor comprising: a second logic circuit 212 that reads out a number of consecutive instruction bundles from the circular buffer 210, and a first logic circuit 213 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 210, and an execution engine 214 for executing the instructions, a circular buffer 215 having a plurality of entries, an instruction cache 216 that stores a cache line of instructions, and a processor comprising: a second logic circuit 217 that reads out a number of consecutive instruction bundles from the circular buffer 210, and a first logic circuit 218 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 210, and an execution engine 219 for executing the instructions, a circular buffer 220 having a plurality of entries, an instruction cache 221 that stores a cache line of instructions, and a processor comprising: a second logic circuit 222 that reads out a number of consecutive instruction bundles from the circular buffer 220, and a first logic circuit 223 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 220, and an execution engine 224 for executing the instructions, a circular buffer 225 having a plurality of entries, an instruction cache 226 that stores a cache line of instructions, and a processor comprising: a second logic circuit 227 that reads out a number of consecutive instruction bundles from the circular buffer 220, and a first logic circuit 228 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 220, and an execution engine 229 for executing the instructions, a circular buffer 230 having a plurality of entries, an instruction cache 231 that stores a cache line of instructions, and a processor comprising: a second logic circuit 232 that reads out a number of consecutive instruction bundles from the circular buffer 230, and a first logic circuit 233 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 230, and an execution engine 234 for executing the instructions, a circular buffer 235 having a plurality of entries, an instruction cache 236 that stores a cache line of instructions, and a processor comprising: a second logic circuit 237 that reads out a number of consecutive instruction bundles from the circular buffer 230, and a first logic circuit 238 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 230, and an execution engine 239 for executing the instructions, a circular buffer 240 having a plurality of entries, an instruction cache 241 that stores a cache line of instructions, and a processor comprising: a second logic circuit 242 that reads out a number of consecutive instruction bundles from the circular buffer 240, and a first logic circuit 243 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 240, and an execution engine 244 for executing the instructions, a circular buffer 245 having a plurality of entries, an instruction cache 246 that stores a cache line of instructions, and a processor comprising: a second logic circuit 247 that reads out a number of consecutive instruction bundles from the circular buffer 240, and a first logic circuit 248 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 240, and an execution engine 249 for executing the instructions, a circular buffer 250 having a plurality of entries, an instruction cache 251 that stores a cache line of instructions, and a processor comprising: a second logic circuit 252 that reads out a number of consecutive instruction bundles from the circular buffer 250, and a first logic circuit 253 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 250, and an execution engine 254 for executing the instructions, a circular buffer 255 having a plurality of entries, an instruction cache 256 that stores a cache line of instructions, and a processor comprising: a second logic circuit 257 that reads out a number of consecutive instruction bundles from the circular buffer 250, and a first logic circuit 258 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 250, and an execution engine 259 for executing the instructions, a circular buffer 260 having a plurality of entries, an instruction cache 261 that stores a cache line of instructions, and a processor comprising: a second logic circuit 262 that reads out a number of consecutive instruction bundles from the circular buffer 260, and a first logic circuit 263 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 260, and an execution engine 264 for executing the instructions, a circular buffer 265 having a plurality of entries, an instruction cache 266 that stores a cache line of instructions, and a processor comprising: a second logic circuit 267 that reads out a number of consecutive instruction bundles from the circular buffer 260, and a first logic circuit 268 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 260, and an execution engine 269 for executing the instructions, a circular buffer 270 having a plurality of entries, an instruction cache 271 that stores a cache line of instructions, and a processor comprising: a second logic circuit 272 that reads out a number of consecutive instruction bundles from the circular buffer 270, and a first logic circuit 273 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 270, and an execution engine 274 for executing the instructions, a circular buffer 275 having a plurality of entries, an instruction cache 276 that stores a cache line of instructions, and a processor comprising: a second logic circuit 277 that reads out a number of consecutive instruction bundles from the circular buffer 270, and a first logic circuit 278 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 270, and an execution engine 279 for executing the instructions, a circular buffer 280 having a plurality of entries, an instruction cache 281 that stores a cache line of instructions, and a processor comprising: a second logic circuit 282 that reads out a number of consecutive instruction bundles from the circular buffer 280, and a first logic circuit 283 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 280, and an execution engine 284 for executing the instructions, a circular buffer 285 having a plurality of entries, an instruction cache 286 that stores a cache line of instructions, and a processor comprising: a second logic circuit 287 that reads out a number of consecutive instruction bundles from the circular buffer 280, and a first logic circuit 288 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 280, and an execution engine 289 for executing the instructions, a circular buffer 290 having a plurality of entries, an instruction cache 291 that stores a cache line of instructions, and a processor comprising: a second logic circuit 292 that reads out a number of consecutive instruction bundles from the circular buffer 290, and a first logic circuit 293 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 290, and an execution engine 294 for executing the instructions, a circular buffer 295 having a plurality of entries, an instruction cache 296 that stores a cache line of instructions, and a processor comprising: a second logic circuit 297 that reads out a number of consecutive instruction bundles from the circular buffer 290, and a first logic circuit 298 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 290, and an execution engine 299 for executing the instructions, a circular buffer 300 having a plurality of entries, an instruction cache 301 that stores a cache line of instructions, and a processor comprising: a second logic circuit 302 that reads out a number of consecutive instruction bundles from the circular buffer 300, and a first logic circuit 303 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 300, and an execution engine 304 for executing the instructions, a circular buffer 305 having a plurality of entries, an instruction cache 306 that stores a cache line of instructions, and a processor comprising: a second logic circuit 307 that reads out a number of consecutive instruction bundles from the circular buffer 300, and a first logic circuit 308 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 300, and an execution engine 309 for executing the instructions, a circular buffer 310 having a plurality of entries, an instruction cache 311 that stores a cache line of instructions, and a processor comprising: a second logic circuit 312 that reads out a number of consecutive instruction bundles from the circular buffer 310, and a first logic circuit 313 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 310, and an execution engine 314 for executing the instructions, a circular buffer 315 having a plurality of entries, an instruction cache 316 that stores a cache line of instructions, and a processor comprising: a second logic circuit 317 that reads out a number of consecutive instruction bundles from the circular buffer 310, and a first logic circuit 318 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 310, and an execution engine 319 for executing the instructions, a circular buffer 320 having a plurality of entries, an instruction cache 321 that stores a cache line of instructions, and a processor comprising: a second logic circuit 322 that reads out a number of consecutive instruction bundles from the circular buffer 320, and a first logic circuit 323 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 320, and an execution engine 324 for executing the instructions, a circular buffer 325 having a plurality of entries, an instruction cache 326 that stores a cache line of instructions, and a processor comprising: a second logic circuit 327 that reads out a number of consecutive instruction bundles from the circular buffer 320, and a first logic circuit 328 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 320, and an execution engine 329 for executing the instructions, a circular buffer 330 having a plurality of entries, an instruction cache 331 that stores a cache line of instructions, and a processor comprising: a second logic circuit 332 that reads out a number of consecutive instruction bundles from the circular buffer 330, and a first logic circuit 333 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 330, and an execution engine 334 for executing the instructions, a circular buffer 335 having a plurality of entries, an instruction cache 336 that stores a cache line of instructions, and a processor comprising: a second logic circuit 337 that reads out a number of consecutive instruction bundles from the circular buffer 330, and a first logic circuit 338 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 330, and an execution engine 339 for executing the instructions, a circular buffer 340 having a plurality of entries, an instruction cache 341 that stores a cache line of instructions, and a processor comprising: a second logic circuit 342 that reads out a number of consecutive instruction bundles from the circular buffer 340, and a first logic circuit 343 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 340, and an execution engine 344 for executing the instructions, a circular buffer 345 having a plurality of entries, an instruction cache 346 that stores a cache line of instructions, and a processor comprising: a second logic circuit 347 that reads out a number of consecutive instruction bundles from the circular buffer 340, and a first logic circuit 348 that divides the cache line into instruction bundles, each of which is written into an entry of the circular buffer 340, and an execution engine 349 for executing the instructions, a circular buffer 350 having a plurality of entries, an instruction cache 351 that stores a cache line of instructions, and a processor comprising: a second logic circuit 352 that reads out a number of consecutive instruction bundles from the circular buffer 350, and a first logic circuit 353 that divides the cache line into instruction bundles, each

	Docum ent ID	U	Title	Current OR
18	US 20030 00969 4 A1	<input checked="" type="checkbox"/>	Hardware architecture, operating system and network transport neutral system, method and computer program product for secure communications and messaging	713/201
19	US 20030 00964 8 A1	<input checked="" type="checkbox"/>	Apparatus for supporting a logically partitioned computer system	711/202
20	US 20030 00526 2 A1	<input checked="" type="checkbox"/>	Mechanism for providing high instruction fetch bandwidth in a multi-threaded processor	712/207
21	US 20020 19909 6 A1	<input checked="" type="checkbox"/>	System and method for secure unidirectional messaging	713/153
22	US 20020 19900 1 A1	<input checked="" type="checkbox"/>	System and method for conducting a secure response communication session	709/227
23	US 20020 19693 5 A1	<input checked="" type="checkbox"/>	Common security protocol structure and mechanism and system and method for using	380/37
24	US 20020 19625 9 A1	<input checked="" type="checkbox"/>	Single semiconductor graphics platform with blending and fog capabilities	345/506
25	US 20020 19450 1 A1	<input checked="" type="checkbox"/>	System and method for conducting a secure interactive communication session	713/201
26	US 20020 19448 3 A1	<input checked="" type="checkbox"/>	System and method for authorization of access to a resource	713/185
27	US 20020 18074 0 A1	<input checked="" type="checkbox"/>	Clipping system and method for a single graphics semiconductor platform	345/506
28	US 20020 17836 0 A1	<input checked="" type="checkbox"/>	System and method for communicating a secure unidirectional response message	713/170
29	US 20020 16591 2 A1	<input checked="" type="checkbox"/>	Secure certificate and system and method for issuing and using same	709/203
30	US 20020 10800 3 A1	<input checked="" type="checkbox"/>	COMMAND QUEUEING ENGINE	710/39
31	US 20020 10551 9 A1	<input checked="" type="checkbox"/>	Clipping system and method for a graphics processing framework embodied on a single semiconductor platform	345/426
32	US 20020 04784 6 A1	<input checked="" type="checkbox"/>	System, method and computer program product for performing a scissor operation in a graphics processing framework embodied on a single semiconductor platform	345/522
33	US 20020 03841 6 A1	<input checked="" type="checkbox"/>	System and method for reading and writing a thread state in a multithreaded central processing unit	712/228
34	US 20020 02755 3 A1	<input checked="" type="checkbox"/>	Diffuse-coloring system and method for a graphics processing framework embodied on a single semiconductor platform	345/426

a circular buffer having N-entries and a read pointer and a write pointer, the circular buffer having a pair of write ports and a single read port, the circular buffer being organized such that every ith entry shares the single read port;

way multiplexer logic for dividing a cache line of the instruction cache into first and second instruction bundles, the first and second instruction bundles being written into consecutive entries of the circular buffer starting from the write pointer;

alignment multiplexer logic for aligning and presenting individual ones of the instruction bundles to the execution core of the processor, wherein j consecutive instruction bundles are read out of the circular buffer at a time starting from the read pointer; and

control logic for selecting instruction bundles to be output from the alignment multiplexer logic.

16. The logic circuit of claim 15 wherein the alignment multiplexer logic comprises:

first and second multiplexers coupled to the single read port of the circular buffer, the first and second multiplexers each outputting an instruction bundle which is respectively latched in the first and second latches prior to dispersal to the execution core.

17. The logic circuit of claim 16 further comprising:

an instruction bypass path coupled between the way multiplexer logic and the alignment multiplexer logic; and

wherein the control logic selects appropriate instruction bundles to be output from the first and second multiplexers depending on a number of instruction bundles consumed by the execution core, and a number of instruction bundles available from the circular buffer.

18. The logic circuit of claim 17 wherein the number of instruction bundles consumed by the execution core in a clock cycle is at most 2.

19. The logic circuit of claim 17 wherein the number of instruction bundles available from the circular buffer in a clock cycle is at most 4.

* * *

3. The processor of claim 1 wherein the number of consecutive instruction bundles is equal to 4.

4. The processor of claim 1 further comprising an instruction bypass path coupled between the first and second logic circuits.

5. The processor of claim 1 wherein the plurality of entries is equal to 8 and the cache line comprises 32 bytes.

6. A processor comprising:

an instruction cache that stores a cache line of instructions;

an execution core for executing the instructions;

a first multiplexer circuit that divides the cache line into a N-entry buffer having a plurality of entries and a single read port, the buffer being organized such that every ith entry shares the single read port, the buffer including a fill pointer that points to a starting entry for filling with an instruction bundle provided by the first logic circuit; and

a second multiplexer circuit that aligns and presents each instruction bundle to the execution core;

drain control logic that selects an appropriate number of instruction bundles read out from the buffer.

7. The processor of claim 6 wherein the cache line comprises 32 bytes.

8. The processor of claim 7 wherein $N=8$ and $i=4$.

9. The processor of claim 6 wherein the number of consecutive instructions is equal to 4.

10. The processor of claim 6 further comprising an instruction bypass path coupled between the first and second multiplexer circuits.

11. The processor of claim 10 wherein the instruction bypass path transfers a pair of instruction bundles.

12. The processor of claim 6 wherein the second multiplexer circuit includes latches that latch the instruction bundles prior to dispersal to the execution core.

13. The processor of claim 10 wherein the instruction bypass path transfers a pair of instruction bundles.

14. The processor of claim 10 wherein the buffer comprises a circular buffer.

15. A logic circuit which supplies instructions from an instruction cache to the execution core of a processor comprising:

	Docum ent ID	U	Title	Current OR
35	US 20020 01386 6 A1	<input checked="" type="checkbox"/>	RETRIEVAL CHANNELS FOR A DATA CONTROLLER	710/22
36	US 20020 00266 7 A1	<input checked="" type="checkbox"/>	System and method for instruction level multithreading in an embedded processor using zero-time context switching	712/228
37	US 20010 01870 1 A1	<input checked="" type="checkbox"/>	Performance enhancements for threaded servers	718/105
38	US 20010 01762 6 A1	<input checked="" type="checkbox"/>	Graphics processing unit with transform module capable of handling scalars and vectors	345/501
39	US 20010 00520 9 A1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for a transform module in a graphics processor	345/506
40	US 66878 12 B1	<input checked="" type="checkbox"/>	Parallel processing apparatus	712/230
41	US 66503 31 B2	<input checked="" type="checkbox"/>	System, method and computer program product for performing a scissor operation in a graphics processing framework embodied on a single semiconductor platform	345/522
42	US 66503 30 B2	<input checked="" type="checkbox"/>	Graphics system and method for processing multiple independent execution threads	345/506
43	US 66503 25 B1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for boustrophedonic rasterization	345/426
44	US 65973 56 B1	<input checked="" type="checkbox"/>	Integrated tessellator in a graphics processing unit	345/423
45	US 65773 09 B2	<input checked="" type="checkbox"/>	System and method for a graphics processing framework embodied utilizing a single semiconductor platform	345/426
46	US 65739 00 B1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for a sequencer in a transform/lighting module capable of processing multiple independent execution threads	345/537
47	US 65642 67 B1	<input checked="" type="checkbox"/>	Network adapter with large frame transfer emulation	709/250
48	US 65156 71 B1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for a vertex attribute buffer in a graphics processor	345/506
49	US 65045 42 B1	<input checked="" type="checkbox"/>	Method, apparatus and article of manufacture for area rasterization using sense points	345/441
50	US 64704 43 B1	<input checked="" type="checkbox"/>	Pipelined multi-thread processor selecting thread instruction in inter-stage buffer based on count information	712/205
51	US 64627 37 B2	<input checked="" type="checkbox"/>	Clipping system and method for a graphics processing framework embodied on a single semiconductor platform	345/426
52	US 64525 95 B1	<input checked="" type="checkbox"/>	Integrated graphics processing unit with antialiasing	345/426
53	US 64496 66 B2	<input checked="" type="checkbox"/>	One retrieval channel in a data controller having staging registers and a next pointer register and programming a context of a direct memory access block	710/23
54	US 64386 71 B1	<input checked="" type="checkbox"/>	Generating partition corresponding real address in partitioned mode supporting system	711/173
55	US 64271 61 B1	<input checked="" type="checkbox"/>	Thread scheduling techniques for multithreaded servers	718/102

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

Page 1 of 1

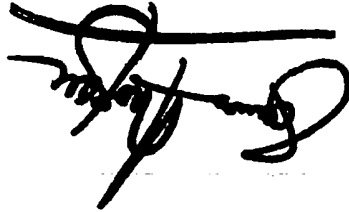
PATENT NO. : 6,539,469 B1
DATED : March 25, 2003
INVENTOR(S) : Jesse Pan

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 1,
Line 30, delete "modern", insert -- modern --.

Signed and Sealed this

Second Day of September, 2003



JAMES E. ROGAN

Director of the United States Patent and Trademark Office

	Docum ent ID	U	Title	Current OR
56	US 64178 51 B1	<input checked="" type="checkbox"/>	Method and apparatus for lighting module in a graphics processor	345/426
57	US 63534 39 B1	<input checked="" type="checkbox"/>	System, method and computer program product for a blending operation in a transform module of a computer graphics pipeline	345/561
58	US 63428 88 B1	<input checked="" type="checkbox"/>	Graphics processing unit with an integrated fog and blending operation	345/426
59	US 63361 50 B1	<input checked="" type="checkbox"/>	Apparatus and method for enhancing data transfer rates using transfer control blocks	710/5
60	US 63245 94 B1	<input checked="" type="checkbox"/>	System for transferring data having a generator for generating a plurality of transfer extend entries in response to a plurality of commands received	710/5
61	US 62232 02 B1	<input checked="" type="checkbox"/>	Virtual machine pooling	718/102
62	US 62125 42 B1	<input checked="" type="checkbox"/>	Method and system for executing a program within a multiscalar processor by processing linked thread descriptors	718/102
63	US 61984 88 B1	<input checked="" type="checkbox"/>	Transform, lighting and rasterization system embodied on a single semiconductor platform	345/426
64	US 60731 59 A	<input checked="" type="checkbox"/>	Thread properties attribute vector based thread selection in multithreading processor	718/103
65	US 59616 39 A	<input checked="" type="checkbox"/>	Processor and method for dynamically inserting auxiliary instructions within an instruction stream during execution	712/242
66	US 59535 20 A	<input checked="" type="checkbox"/>	Address translation buffer for data processing system emulation mode	703/26
67	US 59499 94 A	<input checked="" type="checkbox"/>	Dedicated context-cycling computer with timed context	712/228
68	US 59139 25 A	<input checked="" type="checkbox"/>	Method and system for constructing a program including out-of-order threads and processor and method for executing threads out-of-order	712/206
69	US 58871 66 A	<input checked="" type="checkbox"/>	Method and system for constructing a program including a navigation instruction	718/102
70	US 58128 11 A	<input checked="" type="checkbox"/>	Executing speculative parallel instructions threads with forking and inter-thread communication	712/216
71	US 57817 76 A	<input checked="" type="checkbox"/>	Industrial controller permitting program editing during program execution	717/130
72	US 57428 22 A	<input checked="" type="checkbox"/>	Multithreaded processor which dynamically discriminates a parallel execution and a sequential execution of threads	718/102
73	US 57245 65 A	<input checked="" type="checkbox"/>	Method and system for processing first and second sets of instructions by first and second types of processing systems	712/245
74	US 53576 17 A	<input checked="" type="checkbox"/>	Method and apparatus for substantially concurrent multiple instruction thread processing by a single pipeline processor	712/245
75	US 44424 84 A	<input checked="" type="checkbox"/>	Microprocessor memory management and protection mechanism	711/163

Walk-Up Printing

UserID: f

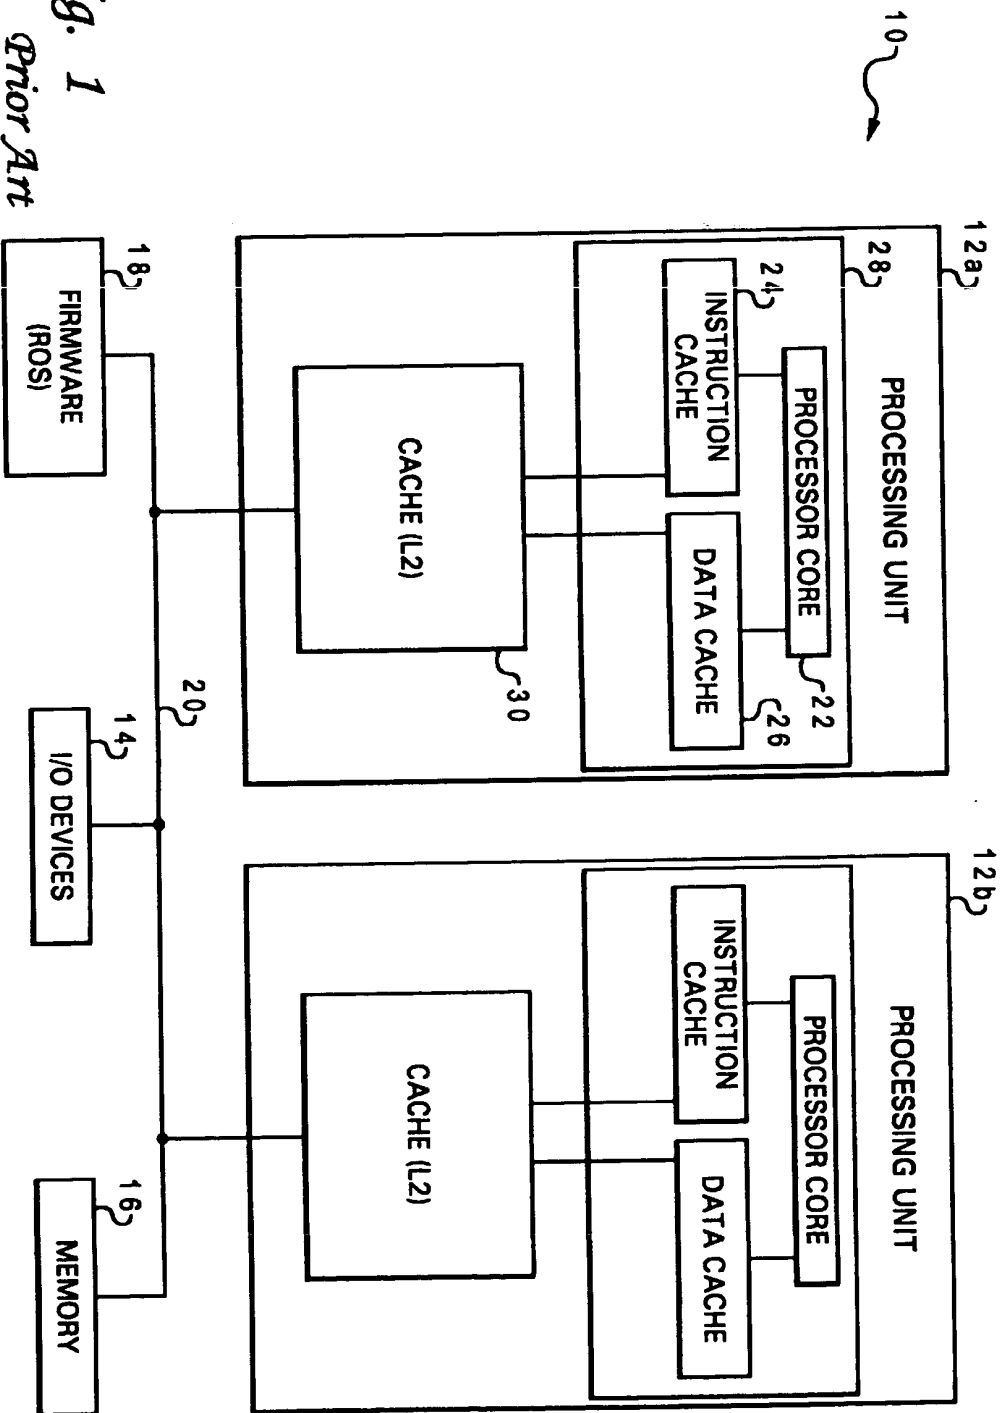
Printer: cpk2_2c21_gbkbpnr

Summary

Document	Pages	Printed	Missed	Copies
US006292845	10	10	0	1
US006324639	53	53	0	1
US006539469	7	7	0	1
US005978896	9	9	0	1
US006065110	9	9	0	1
US20010052053	55	55	0	1
US006691221	19	19	0	1
Total (7)	162	162	0	-

	Docum ent ID	U	Title	Current OR
1	US 20040 01568 4 A1	<input type="checkbox"/>	Method, apparatus and computer program product for scheduling multiple threads for a processor	712/245
2	US 20030 23352 1 A1	<input checked="" type="checkbox"/>	Method and apparatus for converting memory instructions to prefetch operations during a thread switch window	711/137
3	US 20030 23339 4 A1	<input checked="" type="checkbox"/>	Method and apparatus for ensuring fairness and forward progress when executing multiple threads of execution	718/107
4	US 20030 04651 7 A1	<input checked="" type="checkbox"/>	Apparatus to facilitate multithreading in a computer processor pipeline	712/214
5	US 20030 02875 5 A1	<input checked="" type="checkbox"/>	Interprocessor register succession method and device therefor	712/216
6	US 20030 01868 4 A1	<input checked="" type="checkbox"/>	Multi-thread execution method and parallel processor system	718/102
7	US 20030 01447 2 A1	<input checked="" type="checkbox"/>	Thread ending method and device and parallel processor system	718/107
8	US 20030 01447 1 A1	<input checked="" type="checkbox"/>	Multi-thread execution method and parallel processor system	718/107
9	US 20010 05645 6 A1	<input checked="" type="checkbox"/>	PRIORITY BASED SIMULTANEOUS MULTI-THREADING	718/103
10	US 66584 47 B2	<input checked="" type="checkbox"/>	Priority based simultaneous multi-threading	718/103
11	US 63634 53 B1	<input checked="" type="checkbox"/>	Parallel processor with redundancy of processor pairs	711/2
12	US 62405 08 B1	<input checked="" type="checkbox"/>	Decode and execution synchronized pipeline processing using decode generated memory read queue with stop entry to allow execution generated memory read	712/219
13	US 61611 66 A	<input checked="" type="checkbox"/>	Instruction cache for multithreaded processor	711/125
14	US 58729 85 A	<input checked="" type="checkbox"/>	Switching multi-context processor and method overcoming pipeline vacancies	710/1
15	US 58225 78 A	<input checked="" type="checkbox"/>	System for inserting instructions into processor instruction stream in order to perform interrupt processing	712/244
16	US 55533 05 A	<input checked="" type="checkbox"/>	System for synchronizing execution by a processing element of threads within a process using a state indicator	718/106
17	US 55420 58 A	<input checked="" type="checkbox"/>	Pipelined computer with operand context queue to simplify context-dependent execution flow	713/502
18	US 55353 61 A	<input checked="" type="checkbox"/>	Cache block replacement scheme based on directory control bit set/reset and hit/miss basis in a multiheading multiprocessor environment	711/145
19	US 54044 69 A	<input checked="" type="checkbox"/>	Multi-threaded microprocessor architecture utilizing static interleaving	712/215

Fig. 1
Prior Art



	Document ID	U	Title	Current OR
1	DE 10110 504 A1	<input type="checkbox"/>	Use of threads in connection with processor and accessible data by transferring execution control to next thread in queue if first thread is blocked	
2	US 66256 35 B	<input type="checkbox"/>	Threads scheduling method for multi-threaded data processing system, involves allocating prespecified number of instructions to instruction counter which counts execution of instruction by threads sequentially	
3	DE 10110 504 A	<input type="checkbox"/>	Use of threads in connection with processor and accessible data by transferring execution control to next thread in queue if first thread is blocked	

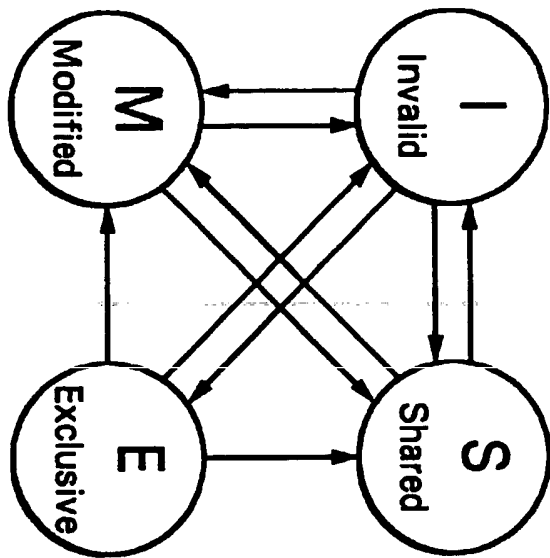


Fig. 2
Prior Art

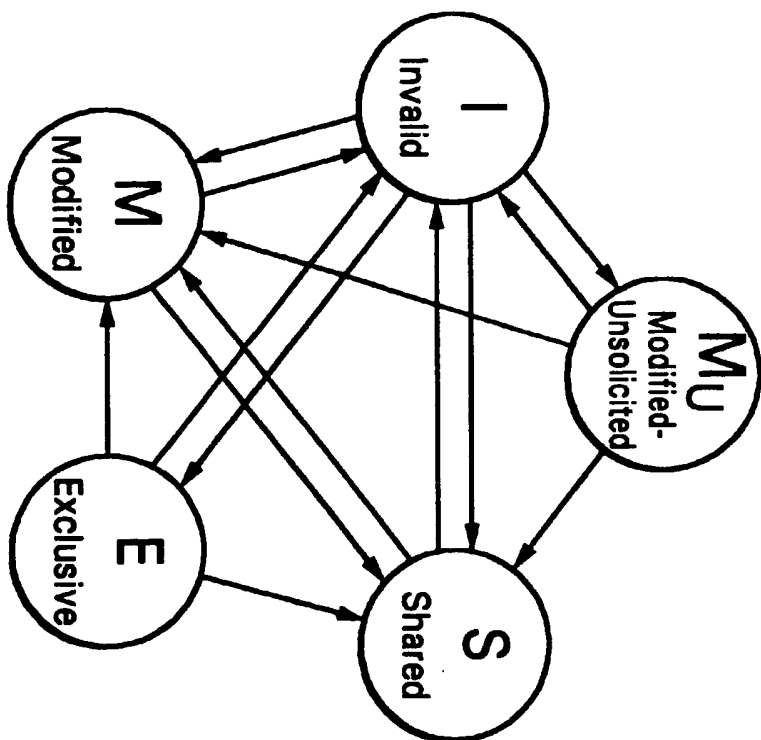


Fig. 3

	L #	Hits	Search Text	DBs
1	L4	493	(select\$3 execut\$3) near20 (thread context) near20 (sequential\$3 robin)	USPAT; US-PGPUB
2	L5	1692	(instruction adj2 (pointer counter) ip (address near10 (fetch\$3 prefetch\$3))) near20 (thread context)	USPAT; US-PGPUB
3	L6	75	4 and 5	USPAT; US-PGPUB
4	L7	34	(select\$3 execut\$3) near20 (thread context) near20 (sequential\$3 robin)	EPO; JPO; DERWENT; IBM_TDB
5	L8	89	(instruction adj2 (pointer counter) ip (address near10 (fetch\$3 prefetch\$3))) near20 (thread context)	EPO; JPO; DERWENT; IBM_TDB
6	L9	1	7 and 8	EPO; JPO; DERWENT; IBM_TDB
7	L10	1889	(instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (thread context)	USPAT; US-PGPUB
8	L11	19	4 and 10 not 6	USPAT; US-PGPUB
9	L12	114	(instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (thread context)	EPO; JPO; DERWENT; IBM_TDB
10	L13	2	7 and "13"	EPO; JPO; DERWENT; IBM_TDB
11	L14	3	7 and 12	EPO; JPO; DERWENT; IBM_TDB

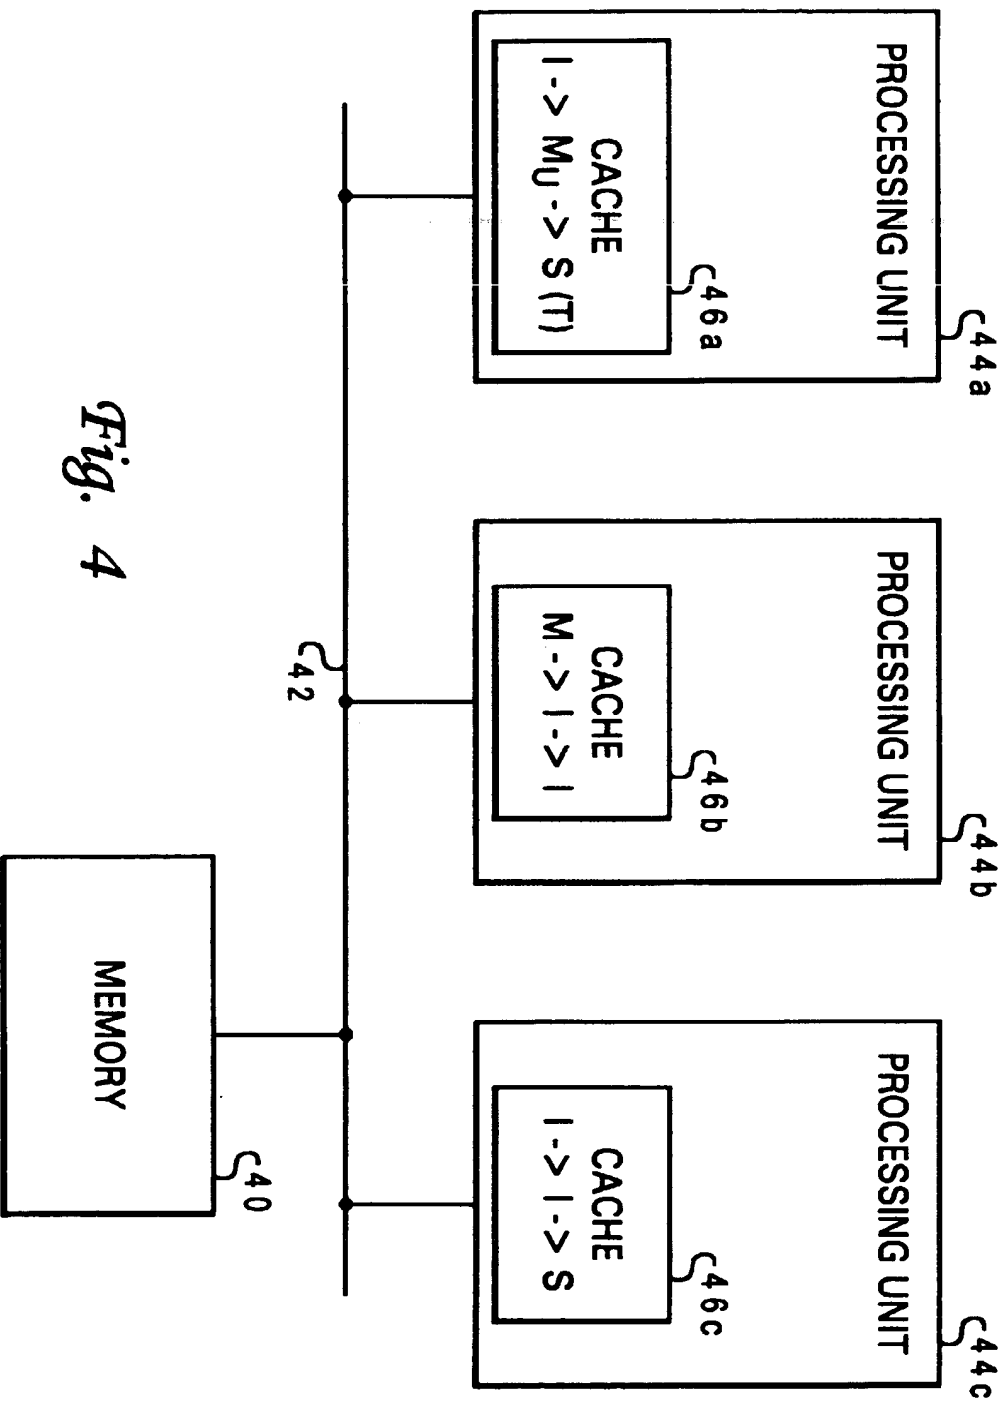


Fig. 4

	Document ID	U	Title	Current OR
1	JP 20032 59440 A	<input type="checkbox"/>	MOBILE PORTABLE TERMINAL AND MANAGEMENT METHOD FOR PRESERVED PDP (PACKET DATA PROTOCOL) CONTEXT USED THEREFOR	
2	JP 20031 74671 A	<input checked="" type="checkbox"/>	MOBILE TERMINAL AND METHOD FOR MANAGING PDP CONTEXT IN ON-STATE AT ALL TIMES	
3	JP 20021 41931 A	<input checked="" type="checkbox"/>	ROUTER AND ROUTE CONTROL METHOD	
4	JP 20020 09938 A	<input checked="" type="checkbox"/>	SYSTEM AND METHOD FOR MULTISERVICE COMMUNICATION	
5	JP 20021 41931 A	<input checked="" type="checkbox"/>	ROUTER AND ROUTE CONTROL METHOD	
6	JP 20020 09938 A	<input checked="" type="checkbox"/>	SYSTEM AND METHOD FOR MULTISERVICE COMMUNICATION	
7	JP 06168 117 A	<input checked="" type="checkbox"/>	DATA PROCESSOR	
8	JP 04288 632 A	<input checked="" type="checkbox"/>	ONE-CHIP MICROCOMPUTER	
9	JP 04097 434 A	<input checked="" type="checkbox"/>	CONTROL SYSTEM FOR PROCESSING ASYNCHRONOUS EVENT	
10	JP 03186 572 A	<input checked="" type="checkbox"/>	THREAD HANDLING METHOD FOR TAIL END	
11	JP 03141 435 A	<input checked="" type="checkbox"/>	PROCESS SWITCHING SYSTEM	
12	JP 40206 6309 A	<input checked="" type="checkbox"/>	TWO-PIECE CLAMPING DEVICE	
13	JP 01292 430 A	<input checked="" type="checkbox"/>	PARALLEL PROCESSING PROCESSOR	
14	WO 31076 16 A1	<input checked="" type="checkbox"/>	METHOD AND APPARATUS FOR INTERNET PROTOCOL HEADERS COMPRESSION INITIALIZATION	
15	WO 30923 14 A1	<input checked="" type="checkbox"/>	OPTIMIZED INFORMATION TRANSFER ASSOCIATED WITH RELOCATION OF AN IP SESSION IN A MOBILE COMMUNICATIONS SYSTEM	
16	WO 30551 70 A1	<input checked="" type="checkbox"/>	METHOD AND SYSTEM FOR SECURE HANDLING OF ELECTRONIC BUSINESS TRANSACTIONS ON THE INTERNET	
17	EP 13220 90 A2	<input checked="" type="checkbox"/>	Mechanism for simplifying roaming in a communications system	
18	WO 20871 89 A1	<input checked="" type="checkbox"/>	SYSTEM FOR EMULATING A TERMINAL WHICH IS ADAPTED TO ACCESS A TELECOMMUNICATION NETWORK FROM A TERMINAL WHICH IS INADAPTED TO SAID ACCESS	
19	WO 20548 11 A1	<input checked="" type="checkbox"/>	POSITIONING OF TERMINAL EQUIPMENT	
20	EP 12138 92 A2	<input checked="" type="checkbox"/>	System and method for implementing a client side HTTP stack	

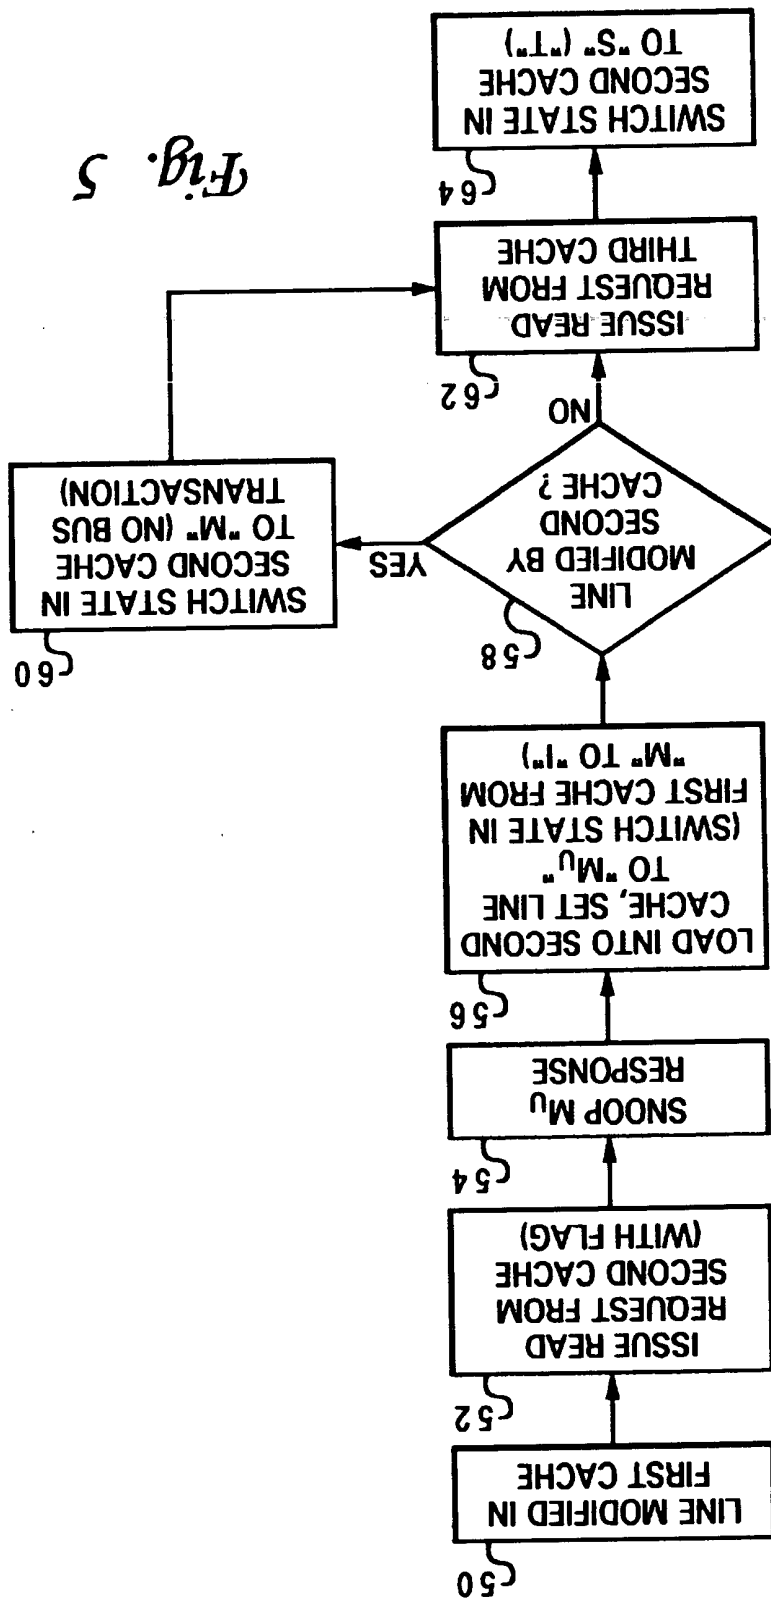


Fig. 5

	Document ID	U	Title	Current OR
21	EP 11331 40 A2	<input checked="" type="checkbox"/>	Multi-service communication system and method	
22	EP 94236 5 A2	<input checked="" type="checkbox"/>	Context controller having instruction-based time slice task switching capability and processor employing the same	
23	WO 97373 03 A1	<input checked="" type="checkbox"/>	TITLE DATA NOT AVAILABLE	
24	EP 58010 9 A2	<input checked="" type="checkbox"/>	Data acces in a RISC digital signal processor.	
25	EP 13438 6 A2	<input checked="" type="checkbox"/>	Method and apparatus for executing object code instructions compiled from a high-level language source.	
26	NNRD4 29148	<input checked="" type="checkbox"/>	Controlled Server Side Execution Environment	
27	NNRD4 2676	<input checked="" type="checkbox"/>	Multiple Terminal Simulation Tool Utilizing Multiple Separate IP Addresses	
28	NN950 541	<input checked="" type="checkbox"/>	Memory Refresh Schemes For Transmission Control Protocol/Internet Protocol	
29	NN941 0425	<input checked="" type="checkbox"/>	Compound Document Client/Server Message Part Handler	
30	NN940 9521	<input checked="" type="checkbox"/>	Fragment Recovery When Implementing a Datagram Fragmentation Scheme	
31	NN930 5211	<input checked="" type="checkbox"/>	Multisequencing a Single Instruction Stream - Concurrent Execution on Alternative Decoder Paths	
32	NB891 0349	<input checked="" type="checkbox"/>	Threaded Code Design for SET Operations	
33	WO 20040 16026 A	<input checked="" type="checkbox"/>	Application level message carrying method for cellular network, involves delivering encapsulated message to receiver application process by transmitting signaling message	
34	US 20040 01311 6 A	<input checked="" type="checkbox"/>	Mobile IP functionality providing method for non mobile IP capable mobile node, involves discovering address of home agent by switching device and advertising address as care-of-address with respect to former address	
35	US 20040 01078 8 A	<input checked="" type="checkbox"/>	Hardware context allocation method for microprocessor, involves permanently allocating hardware content to specific virtual machines based on their resource requirements	
36	WO 20031 07616 A	<input checked="" type="checkbox"/>	Internet protocol headers decompressor node for initializing headers compression, has application module that decompresses application related Internet protocol packets exchanged with one of nodes using decompression context	
37	US 20030 22971 0 A	<input checked="" type="checkbox"/>	Incoming data stream matching method for use in broadband data networking equipment, involves determining potential pattern matches of contexts generated by dividing incoming data stream, in database of known signatures	
38	US 20030 21717 4 A	<input checked="" type="checkbox"/>	Internet protocol session initiating method, involves receiving session request containing host contact and device identifier with device address that is obtained by activation of packet data protocol context	
39	US 66508 77 B	<input checked="" type="checkbox"/>	Physical parameter sensing method for use with radio broadcast context, involves identifying data location for physical parameter transmitted from listener terminal, and returning identified location to listener	
40	WO 20030 92314 A	<input checked="" type="checkbox"/>	Internet protocol session relocation support method, involves determining potential next network node applicability to relocate Internet protocol session by capability information on set of capabilities	
41	CN 14507 50 A	<input checked="" type="checkbox"/>	Method for realizing multi-casting roam	
42	US 66256 35 B	<input checked="" type="checkbox"/>	Threads scheduling method for multi-threaded data processing system, involves allocating prespecified number of instructions to instruction counter which counts execution of instruction by threads sequentially	

CACHE COHERENCY PROTOCOL IN

WHICH A LOAD INSTRUCTION HINT BIT

IS EMPLOYED TO INDICATE

DEALLOCATION OF A MODIFIED CACHE

LINE SUPPLIED BY INTERVENTION

CROSS-REFERENCES TO RELATED

APPLICATIONS

The present invention is related to the following applica-

tions filed concurrently with this application: U.S. patent

application Ser. No. 09/437,179 entitled "HIGH PERFOR-

MANCE MULTIPROCESSOR SYSTEM WITH

MODIFIED-UNSOLICITED CACHE STATE"; U.S. patent

application Ser. No. 09/437,178 entitled "MULTIPROCES-

SOR SYSTEM BUS PROTOCOL WITH COMMAND

AND SNOOP RESPONSES FOR MODIFIED-

UNSOLICITED CACHE STATE"; U.S. patent application

Ser. No. 09/437,177 entitled "MODIFIED-UNSOLICITED

CACHE STATE WITH DYNAMIC HARDWARE/

SOFTWARE CONTROLLABLE MECHANISMS TO OPTI-

MIZE FOR LOAD IMBALANCE"; U.S. patent application

Ser. No. 09/437,180 entitled "PROTOCOL FOR TRANS-

FERRING MODIFIED-UNSOLICITED STATE DURING

DATA INTERVENTION"; U.S. patent application Ser. No.

09/437,181 entitled "CACHE ALLOCATION MECHA-

NISM FOR MODIFIED-UNSOLICITED CACHE

STATES".

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention generally relates to computer systems, and more particularly to a cache coherency pro- col which provides novel mechanisms for handling data in a multi-processor computing system, resulting in improved cache usage and the reduction of bus traffic between hor- zontal caches when a value contained in one cache is modified after being retrieved, in a previously modified form, from a cache of an adjacent processing unit.

2. Description of Related Art

The basic structure of a conventional multi-processor computer system 10 is shown in FIG. 1. Computer system 10 has several processing units, two of which 12a and 12b are depicted, which are connected to various peripheral devices, including input/output (I/O) devices 14 (such as a display monitor, keyboard, graphical pointer (mouse), and a perma- nent storage device or hard disk), memory device 16 (such as random access memory or RAM) that is used by the processing units to carry out program instructions, and firmware 18 whose primary purpose is to seek out and load an operating system from one of the peripherals (usually the permanent memory device) whenever the computer is first turned on. Processing units 12a and 12b communicate with the peripheral devices by various means, including a gen- eralized interconnect or bus 20, or direct memory access channels (not shown). Computer system 10 may have many additional components which are not shown, such as serial, parallel, and universal system bus (USB) ports for connec- tion to, e.g., modems, printers or scanners. There are other components that might be used in conjunction with those shown in the block diagram of FIG. 1; for example, a display adapter might be used to control a video display monitor, a memory controller can be used to access memory 16, etc. The computer can also have more than two processing units. In a symmetric multi-processor (SMP) computer, all of the processing units are generally identical, that is, they all use a common set or subset of instructions and protocols to

operate, and generally have the same architecture. A typical architecture is shown in FIG. 1. A processing unit includes a processor core 22 having a plurality of registers and execution units, which carry out program instructions in order to operate the computer. An exemplary processing unit includes the PowerPC™ processor marketed by Interna- tional Business Machines Corp. The processing unit can also have one or more caches, such as an instruction cache 24 and a data cache 26, which are implemented using high speed memory devices. Caches are commonly used to temporarily store values that might be repeatedly accessed by a processor, in order to speed up processing by avoiding the additional latency of loading the values from memory 16. These caches are referred to as "on-board" when they are integrally packaged with the processor core on a single integrated chip 28. Each cache is associated with a cache controller (not shown) that manages the transfer of data and instructions between the processor core and the cache memory.

A processing unit can include additional caches, such as cache 30, which is referred to as a level 2 (L2) cache since it supports the on-board (level 1) caches 24 and 26. In other words, cache 30 acts as an intermediary between memory 16 and the on-board caches, and can store a much larger amount of information (instructions and data) than the on-board caches can, but at a longer access penalty. For example, cache 30 may be a chip having a storage capacity of 512 kilobytes, while the processor may be an IBM PowerPC™ 604-series processor having on-board caches with 64 kilo- bytes of total storage. Cache 30 is connected to bus 20, and all loading of information from memory 16 into processor core 22 must come through cache 30. Although FIG. 1 depicts only a two-level cache hierarchy, multi-level cache hierarchies can be provided where there are many levels (L3, L4, etc.) of serially connected caches.

In a multi-level cache, if a copy of a value is in every level of the cache, the cache hierarchy is referred to as being "inclusive." It is not necessary, however, to keep a copy of each value in the lower levels, and an inclusivity bit field may be added to the caches to indicate whether or not the cache is inclusive. For example, a three-level cache structure might provide an L3 cache which was not inclusive, such that a value residing in the L2 cache might not be present in the L3 cache. In this example, if an L2 cache issues a read command for a value that is not present in any of the caches of that processing unit, it can be passed to that L2 cache without (necessarily) loading it into the L3 cache.

In an SMP computer, it is important to provide a coherent memory system, that is, to cause write operations to each individual memory location to be serialized in some order for all processors. By way of example, assume a location in memory is modified by a sequence of write operations to take on the values: 1, 2, 3, 4. In a cache coherent system, all processors will observe the writes to a given location to take place in the order shown. However, it is possible for a processing element reading the memory location. A given processing element reading the memory location could see the sequence 1, 3, 4, missing the update to the value 2. A system that implements these properties is said to be "coherent". Nearly all coherence protocols operate only to the granularity of the size of a cache block. That is to say, the coherence protocol controls the movement of and write permissions for operand data or instructions on a cache block basis, and not separately for each individual memory location.

There are a number of protocols and techniques for achieving cache coherence that are known to those skilled in

	Document ID	U	Title	Current OR
43	US 20030 16970 4 A	<input checked="" type="checkbox"/>	Mobile communication terminal, has packet data protocol context management reestablishing context when terminal moves to service area after timer completes measuring full measuring time value before update timer	
44	DE 10205 907 A	<input checked="" type="checkbox"/>	Changing internet protocol access connections to remote IP unit establishes second connection between mobile client and server, using different addresses	
45	EP 13264 08 A	<input checked="" type="checkbox"/>	Enabling several virtual servers to participate in several private network address spaces by applying first IP space ID to translation procedures, enabling selection of current virtual server context	
46	US 20030 12605 9 A	<input checked="" type="checkbox"/>	IP (intellectual property) brokering system for use in anti-terrorism prevention and damage control, has IP brokering unit and communication unit for processing customer requirements and deliverable IP between customers and IP agent	
47	WO 20030 55170 A	<input checked="" type="checkbox"/>	Handling end-to-end Internet business transactions by generating session context with allocated IP address and user ID for validation by transaction manager	
48	KR 20030 54980 A	<input checked="" type="checkbox"/>	Incoming service method in mobile communication packet network	
49	US 20030 11279 4 A	<input checked="" type="checkbox"/>	General packet radio service telecommunications system has port assignment module which sequentially assigns multiple internet protocol addresses to same transmission control protocol ports	
50	US 20030 09754 8 A	<input checked="" type="checkbox"/>	Processing system e.g. pipelined computer processing system, selects address and instructions from respective context registers simultaneously for each cycle of execution of processor unit	
51	US 20030 03896 1 A	<input checked="" type="checkbox"/>	Client message structure for data processing system, includes formats having several message header segments for identifying message, and its property, that are produced in accordance with established protocol	
52	WO 20030 13165 A	<input checked="" type="checkbox"/>	Telecommunications arrangement for mobile packet switched communication system in which support node is interrogated to identify resource wasting hanging PDP contexts	
53	KR 20030 06093 A	<input checked="" type="checkbox"/>	Wap push service method in 3gpp imt-2000 packet network	
54	US 20020 19917 9 A	<input checked="" type="checkbox"/>	Code execution method for copyright protection of patent document, involves selecting entry corresponding to trigger instruction, for executing corresponding auxiliary code	
55	US 20020 15008 2 A	<input checked="" type="checkbox"/>	Voice call processing system for IP telephony, has real time streaming engine with higher priority and normal priority threads for transmitting audio data and silence data to VoIP gateway respectively	
56	KR 20020 77755 A	<input checked="" type="checkbox"/>	Nms platform using multiple thread	
57	WO 20025 4811 A	<input checked="" type="checkbox"/>	Position determination for terminal equipment e.g. mobile telephone in GPRS system that does not use GPS or other costly technologies	
58	US 20020 08784 3 A	<input checked="" type="checkbox"/>	Simultaneous multithreaded processor system delivers instruction pointers for active and inactive threads to processor logic and storage element, respectively	
59	US 20020 07813 5 A	<input checked="" type="checkbox"/>	Data moving method in application layer proxy, involves storing address pointer in send queue of communication connection, so as to move data from one communication connection to another connection	

to implement cache coherency in a system, the processors communicate over a common generalized interconnect (i.e., bus 20). The processors pass messages to read or write the interconnect cache of a processing unit and place in system memory, and the read request from the initiating processor will be satisfied. The scenario just described is commonly referred to as a "snoop-and-push". A read request is snoop-ed on the generalized interconnect which causes the processing unit to "push" the block to the bottom of the hierarchy to satisfy the read request made by the initiating processing unit.

30 memory may have moved from the system memory 16 to
 within caches, the most recent valid copy of a given block or
 memory may be found in the system (as mentioned
 above). If a processor (say 12a) attempts to access a memory
 location not present within its cache hierarchy, the correct
 version of the block, which contains the actual (current)
 value for the memory location, may either be in the system
 memory 16 or in one of more of the caches in another
 processing unit, e.g. processing unit 12b. If the correct
 35 give a "shared" or "not shared" indication for any read they
 (this is accomplished by having the oldest lowest level caches
 other processing unit also has a still active copy of the block
 information allowing the processing unit to determine if any
 operation is not retired, the message usually also includes
 not the read must be retired (i.e., reassessed later). If the read
 30 executes a read it receives a message indicating whether or
 one example of this mechanism, when a processing unit
 generalized interconnect and the inter-cache connections. As
 executes a read it receives a message indicating whether or

For example, consider a processor, say P_2 , attempting to read a location in memory. If first polls its own L_1 cache (24 or 26). If the block is not present in the L_1 cache, the request is forwarded to the L_2 cache (30). If the block is not present in the L_2 cache, the request is forwarded on to lower cache levels, e.g., the L_3 cache. If the block is not present in the lower level caches, the request is then presented on the generalized interconnect (20) to be serviced. Once an operation has been placed on the generalized interconnect, all units in the system because no other processing unit has a block without first communicating with other processing units in the system because no other processing unit has a copy of the block. Therefore, it is possible for a processor to read or write a location without first communicating this intention onto the interconnection, but only where the coherency protocol has ensured that no other processor has an

The foregoing cache coherency technique is implemented in a specific protocol referred to as "MESI," and illustrated in FIG. 2. In this protocol, a cache block can be in one of four states, "M" (Modified), "E" (Exclusive), "S" (Shared) or "I" (Invalid). Under the MESI protocol, each cache entry (e.g., a 32-byte sector) has two additional bits which indicate the state of the entry, out of the four possible states. Depending upon the initial state of the entry and the type of access sought by the requesting processor, the state may be

	Docum ent ID	U	Title	Current OR
60	JP 20021 41931 A	<input checked="" type="checkbox"/>	Router device for internet, has context ID rewriting unit which modifies similar context ID owned by several packets using unique number assigned to each address	
61	WO 20023 2170 A	<input checked="" type="checkbox"/>	Address de-registration in Internet protocol multimedia networks of addresses assigned to users in general packet radio service and Internet protocol multimedia networks for providing application services	
62	WO 20023 2084 A	<input checked="" type="checkbox"/>	Address de-registration from Internet protocol multimedia network to de-register addresses assigned to users using gating network element to detect communication channel deactivation	
63	WO 20021 1397 A	<input checked="" type="checkbox"/>	Header compression context control during handover in mobile data communication networks that prevents errors and communications loss during handover	
64	GB 23627 89 A	<input checked="" type="checkbox"/>	Real-time bit pattern search system in transport packet scheme, e.g. IP, ATM, has tracking unit which tracks search context information for each data stream portion	
65	US 20020 01243 3 A	<input checked="" type="checkbox"/>	Authentication method for a mobile node to a packet data network, e.g. IP, using shared secret key of the mobile and the network authentication center	
66	US 62956 00 B	<input checked="" type="checkbox"/>	Microprocessor for executing multithreaded program in computer, fetches new thread based on stored address in response to cache miss indication signal	
67	US 62955 57 B	<input checked="" type="checkbox"/>	Internet traffic simulator for testing ability of site to handle transport control protocol connections, retrieves information from transport control protocol connection and records statistics related to the information	
68	EP 11331 40 A	<input checked="" type="checkbox"/>	Multi-service communication system has PSTN gateway permitting unidirectional and bidirectional communication between PSTN interface and IP interface	
69	KR 20010 63804 A	<input checked="" type="checkbox"/>	Method for maintaining object consistency in distributed-object system	
70	KR 20010 48669 A	<input checked="" type="checkbox"/>	Method for interfacing application program of mobile event information	
71	KR 20010 35104 A	<input checked="" type="checkbox"/>	Method for registering mobile ip having minimum band in asynchronous mobile communication system	
72	WO 20013 1881 A	<input checked="" type="checkbox"/>	Internet protocol header compression initiation in packet switched network, involves initiating compression, when full header packet with modified header and inserted routing header is received by destination router	
73	KR 20010 07695 A	<input checked="" type="checkbox"/>	Chatting service method and system using virtual chatting server	
74	WO 20010 5171 A	<input checked="" type="checkbox"/>	Access network protocol context management method for use in access network, involves monitoring macro mobility registration at gateway node to determine necessity of access network protocol context	
75	KR 20010 02485 A	<input checked="" type="checkbox"/>	Multi-thread microprocessor for instruction fetch	
76	US 61515 69 A	<input checked="" type="checkbox"/>	User access reestablishing method to programs in computer system involves initiating multitasked execution of identified program with modified context registers of identified program	
77	WO 20006 4203 A	<input checked="" type="checkbox"/>	Multimedia related information transmission involves exchanging information between transport protocol layers in terminal and network device arrangements via octet stream and network transmission protocol layers	

processing. It would be further advantageous if the method could be selectively implemented for those applications which tended to perform such modifications quickly after reading the value from another cache.

SUMMARY OF THE INVENTION

It is therefore one object of the present invention to provide an improved method of maintaining cache coherency in a multi-processor computer system.

It is another object of the present invention to provide such a method that allows a cache to avoid unnecessary bus transactions issued to other horizontal caches.

It is yet another object of the present invention to provide a computer system which makes more efficient use of a cache having a cache line replacement/victimization algorithm.

The foregoing objects are achieved in a method of main-

taining cache coherency, comprising the steps of issuing a load instruction from a first cache of a first processing unit,

and sending a hint bit with the load instruction wherein the hint bit is provided as part of an instruction set architecture

of the computer system and indicates whether a second cache of a second processing unit should deallocate a

corresponding cache line upon sourcing a modified value to the first cache responsive to the load instruction. The modi-

fied value is stored into the cache line of the second cache, prior to said issuing step, and a first coherency state (the

conventional M state) is assigned to the cache line of the second cache indicating that the cache line contains the

modified value. When the hint bit is set to indicate that the second cache should deallocate the cache line, the modified

value is sourced from the cache line of the second cache to a cache line of the first cache, the cache line of the second

cache deallocated, and a second coherency state (the modified-unshared, or M_U state) is assigned to the cache

line of the first cache indicating that the cache line of the first cache contains the value as modified by another processing

unit and that the modified value has not been written to a system memory device.

The above as well as additional objectives, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further

illustrative embodiment thereof, will best be understood by reference to the following detailed description of an

illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

FIG. 1 is a block diagram of a prior-art multi-processor computer system;

FIG. 2 is a state diagram depicting a prior art cache coherency protocol (MESI);

FIG. 3 is a state diagram depicting the cache coherency protocol of the present invention, providing a novel coherency state (M_U) to allow a cache to modify a value, in

response to a store operation, without previously intervening the action to another cache which previously intervened the

requested value;

FIG. 4 is a block diagram illustrating one embodiment of the invention within a multi-processor computing system

having at least three processing units with associated caches, and

A further improvement in accessing cache blocks can be

achieved using the cache coherency protocol. This improvement, referred to as "intervention," allows a cache

having control over a memory block to provide the data in that block directly to another cache requesting the value (for

a read-type operation), in other words, bypassing the need to write the data to system memory and then have the request-

ing processor read it back again from memory. Intervention can generally be performed only by a cache having the value

in a block whose state is Modified or Exclusive. In both of these states, there is only one cache block that has a valid

copy of the value, so it is a simple matter to source (write) the value over the bus 20 without the necessity of first

writing it to system memory. The intervention procedure thus speeds up processing by avoiding the longer process of

writing to and reading from system memory (which actually involves three bus operations and two memory operations).

This procedure not only results in better latency, but also increased bus bandwidth.

There are many variations of the MESI protocol. The tagged ("T") state is used to identify a cache block which is

inconsistent with system memory (i.e., modified) and is further responsible for writing the correct (current) value to

memory upon deallocation (or to pass on the tag to another cache block during intervention). The T state can be used to

share a modified value, by marking one of the sharing blocks as (temporarily) responsible for maintaining a valid copy of

the value. The recently read ("R") state can be used to allow intervention when the value is unmodified but shared among

many caches, so as to conveniently mark a single one of the sharing caches as being responsible for intervention. The

hover ("H") state allows a cache line to maintain an address in the directory even though the corresponding value in the

cache entry array is an invalid copy, so that it can snop the correct value for its processing unit if the value happens to

be broadcast as part of an intervention between the caches of two other processing units.

In many instances, a value that is modified by a store operation, and thus present in a cache of the particular

processing unit that issued the store operation, will subsequently be modified by another processing unit before the

read request and another cache contains a modified value. When one cache issues a

cache line, i.e., switch its state from shared to invalid (I). This follow-up write operation accordingly requires a bus

transaction (e.g., a DClaim instruction) to inform the other cache that it must deallocate the line. This bus transaction to

invalidate a cache line may take as long as 100 processor cycles to complete, which can ultimately delay processing.

It would, therefore, be desirable to devise a method of maintaining cache coherency which allowed a requesting

cache to modify a value provided via intervention in response to a read request, without having to issue a bus

transaction, in order to decrease bus traffic and speed up

	Docum ent ID	U	Title	Current OR
78	JP 20001 12772 A	<input checked="" type="checkbox"/>	Data processing system of digital computer, has decoder which decodes instruction in sequence and detects context call instruction	
79	EP 94236 5 A	<input checked="" type="checkbox"/>	Instruction counter for triggering context controller switch to next background task	
80	US 58898 48 A	<input checked="" type="checkbox"/>	Intelligent peripheral operating method for telecommunication intelligent network	
81	US 57109 23 A	<input checked="" type="checkbox"/>	Communicating active messages among nodes of parallel processing computer system - performing procedure identified by instruction pointer on data structure identified by frame pointer in accordance with micro-thread, so that procedure uses Local Parameters pointer to locate additional parameters associated with procedure	
82	US 57548 30 A	<input checked="" type="checkbox"/>	Web/emulator server for providing persistent connection between client and legacy host system - in which applet code is downloaded to client system in response to receiving uniform resource locator associated with legacy host system	
83	JP 09185 515 A	<input type="checkbox"/>	System timer management method for IPS - by establishing interruption period of system timer using period establishing circuit based on system load condition detected by load detector according to interruption context	
84	JP 09034 848 A	<input type="checkbox"/>	Threading control method for distributed IPS - involves updating memory area containing mutual relationship between threads based on thread which accesses distributed virtual share memory	
85	US 54598 45 A	<input type="checkbox"/>	Instruction pipeline sequencing method - in which state information of instruction travels through pipe stages until instruction execution is completed	
86	EP 58010 9 A	<input type="checkbox"/>	Data processor with disc access in RISC DSP - determines prefetch and poststore addresses according to rule determined by context register with index indicative of this address stored in index register, all together forming a streamer between CPU and data memory	
87	US 54993 49 A	<input type="checkbox"/>	Parallel multi thread data processing system - has instruction pointer and frame pointer which identify instructions to be executed and operated on	
88	EP 13438 6 A	<input type="checkbox"/>	Object code execution appts. from high-level language source - includes host computer with emulator for low-level operations and language processor for high-level	
89	SU 61269 2 A	<input type="checkbox"/>	Conical crusher with slit regulator adjusting ring - has bracket on lock nut carrying lock connected to casing to simplify design	

state to the invalid state. This approach allows the use of the modified-unsolicited state to be optional, as the intervening cache might alternatively send a conventional snoop response, e.g., instructing the requesting cache to assign the shared or tagged state to the cache block. Alternatively, the master (intervening cache) could send a snoop response instructing the requesting cache to simply assign the conventional modified state to the block.

Use of this flag may be programmable via either software or hardware. Software that is visible to the computer's operating system may make a determination that it is desirable to use the M_P state (based on the particular application running), and instruct the operating system to enter an operating mode in which the flag is appropriately set. Hardware may alternatively be provided to dynamically monitor cache usage; for example, the modified-unsolicited state might be disabled if the monitor determines that many cache lines. Each processing unit could have its own monitor, or a system-wide monitor could be provided (e.g., as part of the memory controller). If the flag sent with the request indicates that the M_P state is to be excluded, then conventional intervention occurs.

The programmability feature could further be made visible to a user by adding a hint bit to the load instruction set architecture (ISA) for the processor core (e.g., the ISA of a PowerPC™ 630 processor) could be extended to allow software, via the compiler program, to directly indicate the desirability of using the M_P state for a particular set of load instructions using the hint bit. The availability of the feature as part of the ISA is particularly advantageous since it can be dynamically implemented with each individual instruction.

In an alternative implementation of the present invention, a flag can be sent with the intervention data to indicate that the M_P state may be used. In other words, the requesting cache passes the read operation to the remainder of the memory hierarchy in a conventional fashion, but the intervening cache sends the requested value with a flag indicating that the requesting cache can (should) assign the M_P state to the cache line. This flag thus indicates that the requested value has been supplied via intervention rather than from system memory, and that the line in the intervening cache has already been deallocated. This variation differs from the above-explained implementation since it is not the snooper response that instructs the requesting cache to use the M_P state, but rather the indicating flag is supplied along with the requested data.

In another refinement of the present invention, the eviction/victimization algorithm used by the cache may be modified to further enhance cache efficiency. If a cache miss occurs, and if all of the blocks in a particular congruence class already have valid copies of memory blocks, then one of the cache blocks must be selected for victimization. This selection, in the prior art, is typically performed using an array of bits that indicate which of the blocks is the least-recently used (LRU). For example, there are three LRU bits per block for an 8-way set associative cache. The LRU bits from each block in the class are provided as inputs to a decoder having an 8-bit output to indicate which of the blocks is to be victimized. Several LRU hardware solutions have been implemented to provide "fair" information replacement. Fair replacement, however, often does not provide the best application performance. For some applications, it may be preferable to evict a cache line that is in the M_P state, while for others it may be preferable to

FIG. 5 is a chart illustrating the logic flow according to one implementation of the present invention.

DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

With reference now to FIG. 3, there is depicted a state diagram of one embodiment of the cache coherence protocol of the present invention. This protocol is similar to the prior art MESI protocol of FIG. 2, in that it includes the same four prior states (Modified, Exclusive, Shared and Invalid), but it also includes a new " M_P " state (Modified-*Unsolicited*). The M_P state indicates that: (1) the subject value has been modified (i.e., is not currently consistent with system memory); (2) the value was modified by another processing unit (i.e., not by the processing unit associated with the cache that currently contains the value in the M_P state); and (3) the value is exclusively held (i.e., it is not present in any other horizontally adjacent caches). Because the value is exclusively held, it may be modified in that cache without the necessity of issuing a bus transaction to other horizontal caches in the memory hierarchy.

If a processor core executes a load operation and its cache does not currently contain a valid copy of the requested value, the cache passes the load operation to the remainder of the memory hierarchy to obtain the value. If the value is not present in any cache of the computer system, it is provided by system memory. According to one implementation of the present invention, however, if the value is currently contained in a cache associated with another processing unit, and has already been modified by that processing unit (i.e., is currently in the M state), then the value is retrieved from that cache (via intervention), loaded into a cache of the requesting processor, and assigned the M_P state. The intervening cache immediately deallocates the corresponding cache line. Thus, when the same processor core issues a subsequent store instruction to the same address, the M_P state provides an indication to its cache that the value is not currently contained in any other horizontally adjacent caches within the memory hierarchy. Accordingly, the cache is free to store the modified value, and switch the state of its cache line from M_P to M without issuing any bus transactions, as there is no need to deallocate any cache line(s) in the other cache(s).

In contrast, in the prior art, it would be necessary to issue a deallocation request (e.g., DClaim operation) to the intervening cache which would still contain the value. If the standard MESI protocol were in use, upon intervention from the responding cache, both caches would contain the value in the shared (S) state—the value would also be copied back to system memory during intervention. If the alternative to system memory were implemented, the value would not be written back to system memory, but the responding cache would still hold the value in the S state, and the requesting cache would assign the T state to the cache line. In either case, when the processor core of the requesting cache issued the follow-up write instruction, it would be necessary to deallocate the line held in the S state in the other cache. This deallocation operation would increase bus traffic and delay processing. In a preferred implementation of the present invention, a flag (single-bit field) is sent with the read request to indicate that the requesting cache is capable of utilizing the M_P state (i.e., owning previously modified data). Accordingly, the snooper response from the intervening cache can specifically instruct the requesting cache to utilize that state, in which case the intervening cache will switch from the modified

	Docum ent ID	U	Title	Current OR
1	JP 20022 97414 A	<input type="checkbox"/>	SYSTEM SIMULATOR, SIMULATION METHOD AND SIMULATION PROGRAM	
2	JP 20020 73374 A	<input checked="" type="checkbox"/>	INTERRUPT SIMULATION METHOD AND DEVICE	
3	JP 20022 97414 A	<input checked="" type="checkbox"/>	SYSTEM SIMULATOR, SIMULATION METHOD AND SIMULATION PROGRAM	
4	JP 20020 73374 A	<input checked="" type="checkbox"/>	INTERRUPT SIMULATION METHOD AND DEVICE	
5	JP 09190 348 A	<input checked="" type="checkbox"/>	INSTRUCTION PREFETCH BUFFER CONTROL METHOD AND DEVICE THEREFOR AND INSTRUCTION PREFETCH BUFFER FLUSH METHOD	
6	JP 08227 364 A	<input checked="" type="checkbox"/>	METHOD AND DEVICE FOR EXECUTING SEQUENTIAL MULTITHREAD	
7	JP 03137 702 A	<input checked="" type="checkbox"/>	NUMERICALLY CONTROLLED THREAD CUTTING DEVICE	
8	DE 10110 504 A1	<input checked="" type="checkbox"/>	Use of threads in connection with processor and accessible data by transferring execution control to next thread in queue if first thread is blocked	
9	WO 97457 95 A1	<input checked="" type="checkbox"/>	PARALLEL PROCESSOR WITH REDUNDANCY OF PROCESSOR PAIRS	
10	WO 97076 21 A1	<input checked="" type="checkbox"/>	PARALLEL EXECUTION OF REQUESTS IN OSI AGENTS	
11	EP 73547 5 A2	<input checked="" type="checkbox"/>	Method and apparatus for managing objects in a distributed object operating environment	
12	EP 61736 1 A2	<input checked="" type="checkbox"/>	Scheduling method and apparatus for a communication network.	
13	US 20040 00302 3 A	<input checked="" type="checkbox"/>	Computer system processor selection method for loading processing thread, involves overwriting candidate/volunteer processor information based on comparison of load between candidate/volunteer processor	
14	US 66256 35 B	<input checked="" type="checkbox"/>	Threads scheduling method for multi-threaded data processing system, involves allocating prespecified number of instructions to instruction counter which counts execution of instruction by threads sequentially	
15	US 20030 04651 7 A	<input checked="" type="checkbox"/>	Multithreading apparatus for computer processor pipeline, has control unit statically scheduled to execute multiple threads in round robin succession to eliminate need for communication between pipeline stages	
16	JP 20030 52687 A	<input checked="" type="checkbox"/>	Multislice/cone beam computerized tomography devices undergo several interrelated image data corrective procedures of varying relevance/emphasis	
17	US 20020 10800 3 A	<input checked="" type="checkbox"/>	Data controller for peripheral device, generates transfer extend entries/retrieval for data transfer, automatically	
18	WO 20025 4246 A	<input checked="" type="checkbox"/>	Method for testing and monitoring parallel sequentially executed process sections, or threads, in an existing software development system by extension of thread handlers to include event tracers	
19	US 63341 71 B	<input checked="" type="checkbox"/>	Write-combining device for uncacheable stores in computer system, receives two stores for write-combining, if both stores are sequentially executed from same thread	

compiling a program having a plurality of instructions including at least one load instruction which provides a hint bit to indicate whether a cache should deallocate a corresponding cache line upon sourcing the requested value by intervention to a cache of another processing unit;

in response to execution of the load instruction, issuing a read request with the hint bit from a first cache of a first processing unit to a second cache of a second processing unit;

in response to receiving the requested value in a cache line of the first cache assigning a first coherency state to the cache line of the first cache indicating that the cache contains the value;

3. The method of claim 2 further comprising the steps of: setting the hint bit to indicate that the second cache should deallocate the cache line;

sourcing the modified value from the cache line of the second cache to a cache line of the first cache; and deallocating the cache line of the second cache.

4. The method of claim 3 further comprising the step of sending a snoop response from the second cache instructing the first cache to utilize the first coherency state.

5. The method of claim 3 further comprising the steps of: issuing a store operation for an address associated with the modified value from the first processing unit, after said sourcing step; and

modifying the value in the cache line of the first cache in response to the store operation, wherein said modifying step occurs without issuing any bus transaction from the first cache.

6. The method of claim 5 further comprising the step of assigning the second coherency state to the cache line of the first cache in response to said modifying step.

7. The method of claim 3 further comprising the step of modifying a plurality of victimization priority bits associated with the cache line of the first cache in response to said step of assigning the first coherency state.

8. A computer system comprising: a system memory device; a plurality of processing units each having a cache, each of said processing units adapted to execute a load instruction which provides a hint bit to indicate whether a cache containing a requested value in modified form should deallocate a cache line containing the requested value upon sourcing the requested value by intervention to a cache of another processing unit;

bus means for interconnecting said processing units and said system memory device; and

in response to execution of the load instruction, issuing a read request with the hint bit from a first cache of a first processing unit to a second cache of a second processing unit and, responsive to receiving the requested value in a cache line of the first cache, for assigning a first coherency state to the first cache line of said first cache indicating that said

never victimize M_V data, if possible. Accordingly, the LRU bits may be set for M_V data depending upon the particular application. For example, when a value is loaded into a prior art cache, its LRU bits would be set to indicate that it is the most-recently used (MRU) block, but the present invention can override this setting for M_V data, to instead set the bits to indicate that the block is the least-recently used. Conversely, an LRU decoder can be designed such that M_V data is never evicted (unless all blocks in the congruence class are in the M_V state).

FIG. 4 illustrates one embodiment of a multi-processor computer system in which the present invention can be practiced. The invention could be applied to computer systems having new hardware components not shown in FIG. 4, or having other interconnection architectures, so those skilled in the art will appreciate that the present invention is not limited to the generalized system shown in that figure. In this embodiment, the hierarchy includes a system memory device 40 connected via a system bus 42 to a plurality of processing units, 44a, 44b and 44c. Each of these processing units has at least one cache 46a, 46b and 46c (the invention could be applied in multi-level cache hierarchies as well). As seen in FIG. 4, when a cache line is read from cache 46b which already has the line in the modified (M) state, the line goes from invalid (I) to the modified-unshared (MU) in the requesting cache 46a, and the corresponding line in intervening cache 46b is immediately deallocated (goes from M to I). Thereafter, if a store operation were to hit cache 46a, the line would go from modified-unshared to modified (M) (not illustrated in FIG. 4), without any bus transactions between cache 46a and cache 46b. If, however, no store operation occurs prior to a load request from cache 46c then, as illustrated in FIG. 4, cache 46a can intervene the value and switch its state to shared (S), with the value being written to system memory 40 as well. The line in cache 46c is also assigned the shared state. As indicated by the parentheses in FIG. 4, the line in cache 46a could switch to the tagged (T) state if the value is not to be immediately written back to system memory 40.

The logic flow for this implementation is shown in FIG. 5. First, a line is modified in one of the caches (50). A read request is then issued from a second cache, which may be issued with the flag discussed above (52). A snoop response is formulated to indicate that the line can be loaded in the M_V state in the second cache (54). The cache line is then loaded into the second cache, and set to the M_V state (56); the line in the first cache switches from the M state to the I state. Thereafter, if the line is modified (58), then the second cache can switch to the M state with no bus transaction (60). Whether the line is modified or not, if a read request is issued from a third cache (62), the line in the second cache switches to the S (or T) state (64) when it intervenes the data to the third cache.

Although the invention has been described with reference to specific embodiments, this description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiments, as well as alternative embodiments, will become apparent to persons skilled in the art upon reference to the description of the invention. It is therefore contemplated that such modifications can be made without departing from the spirit or scope of the present invention as defined in the appended claims.

What is claimed is:

1. A method of maintaining cache coherency in a computer system having a plurality of processing units, each processing unit having at least one cache, comprising the steps of:

FIG. 4 illustrates one embodiment of a multi-processor computer system in which the present invention can be practiced. The invention could be applied to computer systems having new hardware components not shown in FIG. 4, or having other interconnection architectures, so those skilled in the art will appreciate that the present invention is not limited to the generalized system shown in that figure. In this embodiment, the hierarchy includes a system memory device 40 connected via a system bus 42 to a plurality of processing units, 44a, 44b and 44c. Each of these processing units has at least one cache 46a, 46b and 46c (the invention could be applied in multi-level cache hierarchies as well). As seen in FIG. 4, when a cache line is read from cache 46b which already has the line in the modified (M) state, the line goes from invalid (I) to the modified-unshared (MU) in the requesting cache 46a, and the corresponding line in intervening cache 46b is immediately deallocated (goes from M to I). Thereafter, if a store operation were to hit cache 46a, the line would go from modified-unshared to modified (M) (not illustrated in FIG. 4), without any bus transactions between cache 46a and cache 46b. If, however, no store operation occurs prior to a load request from cache 46c then, as illustrated in FIG. 4, cache 46a can intervene the value and switch its state to shared (S), with the value being written to system memory 40 as well. The line in cache 46c is also assigned the shared state. As indicated by the parentheses in FIG. 4, the line in cache 46a could switch to the tagged (T) state if the value is not to be immediately written back to system memory 40.

The logic flow for this implementation is shown in FIG. 5. First, a line is modified in one of the caches (50). A read request is then issued from a second cache, which may be issued with the flag discussed above (52). A snoop response is formulated to indicate that the line can be loaded in the M_V state in the second cache (54). The cache line is then loaded into the second cache, and set to the M_V state (56); the line in the first cache switches from the M state to the I state. Thereafter, if the line is modified (58), then the second cache can switch to the M state with no bus transaction (60). Whether the line is modified or not, if a read request is issued from a third cache (62), the line in the second cache switches to the S (or T) state (64) when it intervenes the data to the third cache.

Although the invention has been described with reference to specific embodiments, this description is not meant to be construed in a limiting sense. Various modifications of the disclosed embodiments, as well as alternative embodiments, will become apparent to persons skilled in the art upon reference to the description of the invention. It is therefore contemplated that such modifications can be made without departing from the spirit or scope of the present invention as defined in the appended claims.

What is claimed is:

1. A method of maintaining cache coherency in a computer system having a plurality of processing units, each processing unit having at least one cache, comprising the steps of:

FIG. 4 illustrates one embodiment of a multi-processor computer system in which the present invention can be practiced. The invention could be applied to computer systems having new hardware components not shown in FIG. 4, or having other interconnection architectures, so those skilled in the art will appreciate that the present invention is not limited to the generalized system shown in that figure. In this embodiment, the hierarchy includes a system memory device 40 connected via a system bus 42 to a plurality of processing units, 44a, 44b and 44c. Each of these processing units has at least one cache 46a, 46b and 46c (the invention could be applied in multi-level cache hierarchies as well). As seen in FIG. 4, when a cache line is read from cache 46b which already has the line in the modified (M) state, the line goes from invalid (I) to the modified-unshared (MU) in the requesting cache 46a, and the corresponding line in intervening cache 46b is immediately deallocated (goes from M to I). Thereafter, if a store operation were to hit cache 46a, the line would go from modified-unshared to modified (M) (not illustrated in FIG. 4), without any bus transactions between cache 46a and cache 46b. If, however, no store operation occurs prior to a load request from cache 46c then, as illustrated in FIG. 4, cache 46a can intervene the value and switch its state to shared (S), with the value being written to system memory 40 as well. The line in cache 46c is also assigned the shared state. As indicated by the parentheses in FIG. 4, the line in cache 46a could switch to the tagged (T) state if the value is not to be immediately written back to system memory 40.

	Docum ent ID	U	Title	Current OR
20	WO 20018 8713 A	<input checked="" type="checkbox"/>	Garbage collection process for computer memory pushes identifiers onto thread work queue and pops task identifiers from thread queue ends	
21	DE 10110 504 A	<input checked="" type="checkbox"/>	Use of threads in connection with processor and accessible data by transferring execution control to next thread in queue if first thread is blocked	
22	US 20020 00266 7 A	<input checked="" type="checkbox"/>	Embedded processor architecture for enabling multithreading, invokes zero-time context switches between end and beginning of program instruction execution states in threads	
23	US 62431 07 B	<input checked="" type="checkbox"/>	Graphics processor system performance optimizing method involves moving master thread between slave threads of respective processors to cause each processor to execute its graphics pipeline	
24	US 62054 94 B	<input checked="" type="checkbox"/>	Interface controller for disk drive, includes command queuing engine to generate data transfer descriptor for each initiator command and to form thread of sequential data transfers	
25	RD 43314 9 A	<input checked="" type="checkbox"/>	Improving performance score of multi-thread sequential read	
26	JP 10011 301 A	<input checked="" type="checkbox"/>	Multitask processing apparatus used in microprocessor - has memory for context provided separately from internal bus to store each context by which tasks are to be performed sequentially	
27	US 56969 15 A	<input checked="" type="checkbox"/>	Computer controlled display system associated routes user action controlling - determining if computer controlled display system is in first context then automatically sequentially executes each routine of first number of routines	
28	EP 90165 9 B	<input checked="" type="checkbox"/>	Parallel processor with redundancy of processor pairs - with direct pairing between processors of separate memory buses, such that two tightly coupled processors can reciprocally synchronise themselves and share internal register files	
29	US 58988 32 A	<input checked="" type="checkbox"/>	Organising CMIP request execution in OSI environment - providing sub threads within each main thread for simultaneous processing of multiple CMIP requests	
30	EP 83935 0 B	<input checked="" type="checkbox"/>	Transaction synchronisation method for processing system having agents and coordinator - sending vote indicating availability to commit from each agent to coordinator, and determining commit or back-out decision by coordinator when all votes are received	
31	US 54780 51 A	<input checked="" type="checkbox"/>	Mfg. elongated belt having sequentially located plastic components - by moulding component part and guide means with location determining means interconnected by cross-piece onto thread of compatible material at moulding station.	
32	EP 64908 7 A	<input checked="" type="checkbox"/>	Computer system for software function selection - has filtering mechanism for finding and selecting collection of software functions useful for user from collection of available software provided by software developers	
33	EP 64448 6 A	<input checked="" type="checkbox"/>	Management system for data access in computer disc memory system - groups together large data access tasks against partitions onto physical disc to be read from or written to, and executes tasks serially in accordance with disc task list	
34	US 53902 81 A	<input type="checkbox"/>	Deducing user intent and providing computer implemented services for pen-based computer - comparing new events with events stored in database, collecting into set, comparing to intent templates, and selecting best-guess hypothesis	

11

cache line of said first cache contains the value as modified by another processing unit and that the modified value has not been written to said system memory.

9. The computer system of claim 8 wherein said cache coherency means stores the modified value into said cache line of said second cache, prior to said issuing, and assigns a second coherency state to said cache line of said second cache indicating that said cache line contains the modified value.

10. The computer system of claim 9 wherein said cache coherency means further sets the hint bit to indicate that the second cache should deallocate the cache line, sources the modified value from said cache line of said second cache to a cache line of said first cache, and deallocates said cache line of said second cache.

11. The computer system of claim 10 wherein said cache coherency means further sends a snoop response from said

12

second cache instructing said first cache to utilize the first coherency state.

12. The computer system of claim 10 wherein said cache coherency means further modifies the value in said cache line of the first cache in response to a store operation for an address associated with the value issued by said first processing unit, without issuing any bus transaction from said first cache.

13. The computer system of claim 12 wherein said cache coherency means further assigns the second coherency state to said cache line of said first cache in response to modifying the value in said cache line of said first cache.

14. The computer system of claim 10 wherein said cache coherency means further modifies a plurality of victimization priority bits associated with said cache line of said first cache in response to said assigning of the first coherency state to said cache line of said first cache.

* * *

	L #	Hits	Search Text	DBs
1	L1	83	(plural plurality multiple multiplicity several number) adj5 ((instruction adj2 (pointer counter) ip ((instruction address) near10 (fetch\$3 prefetch\$3))) near20 (multithread\$3 thread context)	USPAT; US-PGPUB

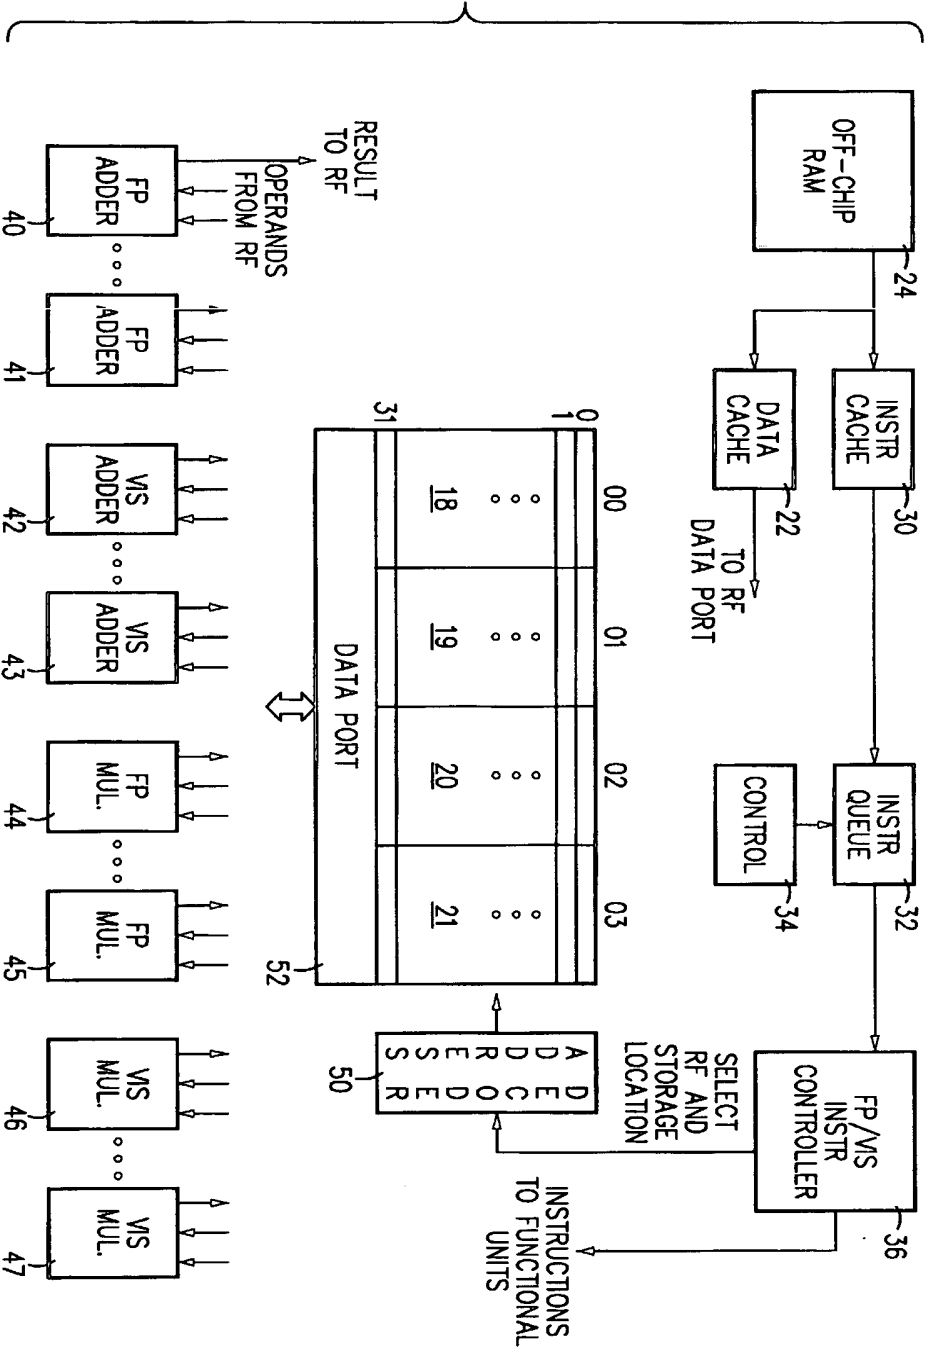


FIG. 6

	Docum ent ID	U	Title	Current OR
1	US 20040 06871 7 A1	<input type="checkbox"/>	System and method of utilizing a hardware component to execute an interpretive language	717/142
2	US 20040 06224 5 A1	<input checked="" type="checkbox"/>	TCP/IP offload device	370/392
3	US 20040 05452 2 A1	<input checked="" type="checkbox"/>	System and method to access web resources from wireless devices	704/8
4	US 20040 04960 0 A1	<input checked="" type="checkbox"/>	Memory management offload for RDMA enabled network adapters	709/250
5	US 20040 04736 6 A1	<input checked="" type="checkbox"/>	Method for dynamic flow mapping in a wireless network	370/466
6	US 20040 01989 1 A1	<input checked="" type="checkbox"/>	Method and apparatus for optimizing performance in a multi-processing system	718/102
7	US 20040 01588 8 A1	<input checked="" type="checkbox"/>	Processor system including dynamic translation facility, binary translation program that runs in computer having processor system implemented therein, and semiconductor device having processor system implemented therein	717/136
8	US 20040 01067 9 A1	<input checked="" type="checkbox"/>	Reducing processor energy consumption by controlling processor resources	713/1
9	US 20030 21717 4 A1	<input checked="" type="checkbox"/>	Establishing an IP session between a host using SIP and a device without an IP address	709/237
10	US 20030 15430 7 A1	<input checked="" type="checkbox"/>	Method and apparatus for aggregate network address routes	709/245
11	US 20030 11279 4 A1	<input checked="" type="checkbox"/>	System and method for multiple PDP contexts with a single PDP address at a GGSN	370/352
12	US 20030 09754 7 A1	<input checked="" type="checkbox"/>	Context scheduling	712/228
13	US 20030 07184 7 A1	<input checked="" type="checkbox"/>	Notification of messages to a terminal by means of vector image	345/772
14	US 20030 06021 0 A1	<input checked="" type="checkbox"/>	System and method for providing real-time and non-real-time services over a communications system	455/452 .1
15	US 20030 02623 0 A1	<input checked="" type="checkbox"/>	Proxy duplicate address detection for dynamic address allocation	370/338
16	US 20030 01447 3 A1	<input checked="" type="checkbox"/>	Multi-thread executing method and parallel processing system	718/107
17	US 20030 00514 4 A1	<input checked="" type="checkbox"/>	EFFICIENT CLASSIFICATION MANIPULATION AND CONTROL OF NETWORK TRANSMISSIONS BY ASSOCIATING NETWORK FLOWS WITH RULE BASED FUNCTIONS	709/235

g

Printed by HPS Server
for

Walk-Up_Printing

Printer: cpk2_2c21_gbkbptr

Date: 04/07/04

Time: 09:19:27

Document Listing

Document	Selected Pages	Page Range	Copies
US006374333	11	1 - 11	1
US006389512	14	1 - 14	1
US006405307	18	1 - 18	1
US006536034	16	1 - 16	1
Total (4)	59	-	-

	Docum ent ID	U	Title	Current OR
18	US 20020 19161 2 A1	<input checked="" type="checkbox"/>	Method and apparatus for automatically determining an appropriate transmission method in a network	370/392
19	US 20020 18778 8 A1	<input checked="" type="checkbox"/>	Mobile communications system	455/450
20	US 20020 18144 8 A1	<input checked="" type="checkbox"/>	Prevention of spoofing in telecommunications systems	370/352
21	US 20020 13620 6 A1	<input checked="" type="checkbox"/>	Recursive query for communications network data	370/352
22	US 20020 10796 4 A1	<input checked="" type="checkbox"/>	Mobile communication system and data transferring method for use with mobile communication system	709/225
23	US 20020 09984 4 A1	<input checked="" type="checkbox"/>	Load balancing and dynamic control of multiple data streams in a network	709/232
24	US 20020 06243 5 A1	<input checked="" type="checkbox"/>	PRIORITIZED INSTRUCTION SCHEDULING FOR MULTI-STREAMING PROCESSORS	712/7
25	US 20020 01071 0 A1	<input checked="" type="checkbox"/>	Method for characterizing a complex system	715/500
26	US 20010 03744 8 A1	<input checked="" type="checkbox"/>	Input replicator for interrupts in a simultaneous and redundantly threaded processor	712/244
27	US 20010 03744 7 A1	<input checked="" type="checkbox"/>	Simultaneous and redundantly threaded processor branch outcome queue	712/239
28	US 20010 03625 5 A1	<input checked="" type="checkbox"/>	Methods and apparatus for providing speech recognition services to communication system users	379/88. 01
29	US 20010 02947 8 A1	<input checked="" type="checkbox"/>	System and method for supporting online auctions	705/37
30	US 67213 33 B1	<input checked="" type="checkbox"/>	Point to point protocol multiplexing/demultiplexing method and apparatus	370/469
31	US 67078 13 B1	<input checked="" type="checkbox"/>	Method of call control to minimize delays in launching multimedia or voice calls in a packet-switched radio telecommunications network	370/356
32	US 65325 54 B1	<input checked="" type="checkbox"/>	Network event correlation system using formally specified models of protocol behavior	714/43
33	US 65196 36 B2	<input checked="" type="checkbox"/>	Efficient classification, manipulation, and control of network transmissions by associating network flows with rule based functions	709/223
34	US H0020 51 H	<input checked="" type="checkbox"/>	System and method for providing multiple quality of service classes	370/395 .21
35	US 64775 62 B2	<input checked="" type="checkbox"/>	Prioritized instruction scheduling for multi-streaming processors	718/108
36	US 64346 76 B1	<input checked="" type="checkbox"/>	FIFO with random re-read support and its application	711/154



Patent Number:

Prabhun et al.

[54] AUXILIARY REGISTER FILE ACCESSING TECHNIQUE

AUXILIARY REGISTER FILE ACCESSING TECHNIQUE

Ferrellto; Eric T. Anderson, both of Sunnyvale; James A. Bauman, San Jose, all of Calif.

Jose, all of Calif.
Sun Microsystems, Inc., Palo Alto, Calif.

[21] Appl. No.: 787,339

[22] Filed: Jan. 27, 1997

[51] Int. Cl.⁶ G06F 12/02

U.S. CI. 364/DIG 1 MS Etc.

364/DIG. 2 MS File; 395/376; 711/2, 5,

01/01/2007 10:01:01

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,070,703 1/1978 Neg. 711/5

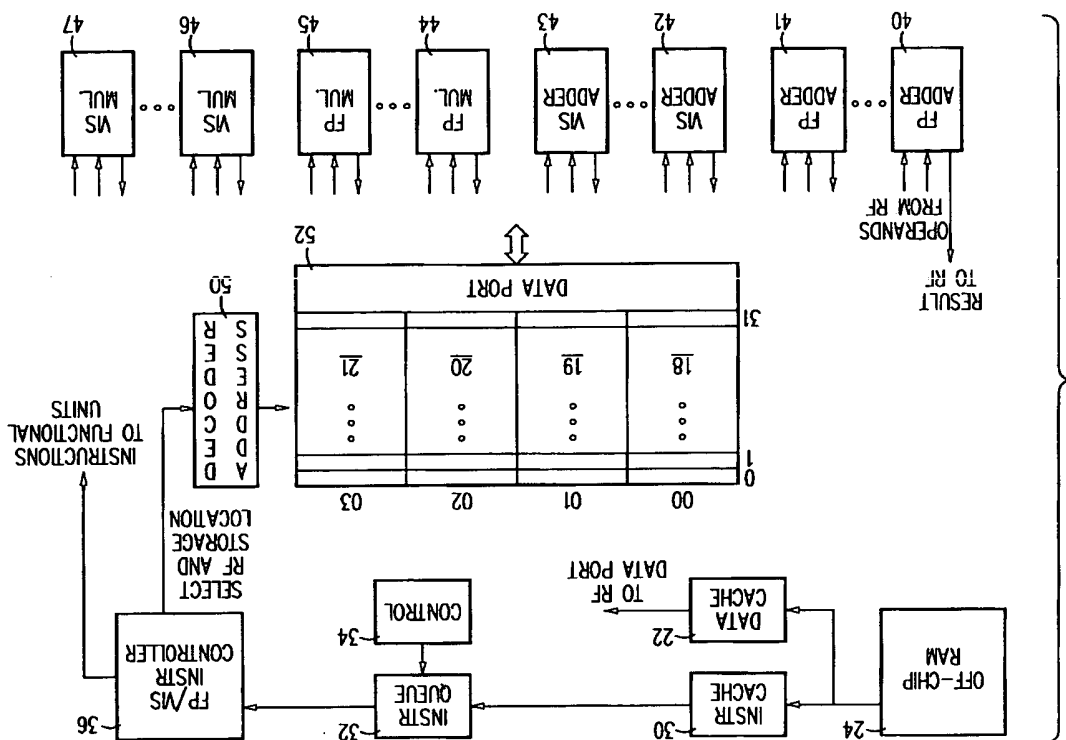
16 Claims, 4 Drawing Sheets

Certain bits in existing op code formats for a processor do not change from one instruction to another when particular classes of instructions are used. Applications optionally utilize one or more of these bits to identify one of a plurality of different register files from which to retrieve operands or to store the results of an operation. These bits along with allocated address bits in predetermined address fields now allow the processor to address many more registers. This can be used to increase the performance of the processor. Those programs not utilizing the bits outside of the address fields for designating a particular register file are backwards compatible with the modified processor.

ABSTRACT

*Primary Examiner—Robert B. Harrell
Attorney, Agent, or Firm—The Gunnison Law Firm*

4,318,175	3/1982	Hawley	7/11/5
5,615,348	3/1997	Koyno et al.	7/11/5



	Docum ent ID	U	Title	Current OR
37	US 63380 78 B1	<input checked="" type="checkbox"/>	System and method for sequencing packets for multiprocessor parallelization in a computer network system	718/102
38	US 62956 00 B1	<input checked="" type="checkbox"/>	Thread switch on blocked load or store using instruction thread field	712/228
39	US 62955 57 B1	<input checked="" type="checkbox"/>	Apparatus for simulating internet traffic	709/224
40	US 62928 88 B1	<input checked="" type="checkbox"/>	Register transfer unit for electronic processor	712/225
41	US 62721 48 B1	<input checked="" type="checkbox"/>	Scheme for reliable communications via radio and wire networks using transport layer connection	370/469
42	US 62333 15 B1	<input checked="" type="checkbox"/>	Methods and apparatus for increasing the utility and interoperability of peripheral devices in communications systems	379/88. 01
43	US 62298 80 B1	<input checked="" type="checkbox"/>	Methods and apparatus for efficiently providing a communication system with speech recognition capabilities	379/88. 01
44	US 61547 77 A	<input checked="" type="checkbox"/>	System for context-dependent name resolution	709/227
45	US 58549 13 A	<input type="checkbox"/>	Microprocessor with an architecture mode control capable of supporting extensions of two distinct instruction-set architectures	712/210

10
A

	rd	op3	rs1	14 13	19 18	25 24	31 30 29
rs2	opf			5 4	0		

FIG. 1 (PRIOR ART)

Opcode	opf	Operation
FPADD16	0 0101 0000	FOUR 16-BIT ADD
FPADD16S	0 0101 0001	TWO 16-BIT ADD
FPADD32	0 0101 0010	TWO 32-BIT ADD
FPADD32S	0 0101 0011	ONE 32-BIT ADD
FPSUB16	0 0101 0100	FOUR 16-BIT SUBTRACT
FPSUB16S	0 0101 0101	TWO 16-BIT SUBTRACT
FPSUB32	0 0101 0110	TWO 32-BIT SUBTRACT
FPSUB32S	0 0101 0111	ONE 32-BIT SUBTRACT

FIG. 2 (PRIOR ART)

Opcode	op3	opf	Operation
FMLS	11 0100	0 0100 1001	MULTIPLY SINGLE
FMULD	11 0100	0 0100 1010	MULTIPLY DOUBLE
FMULq	11 0100	0 0100 1011	MULTIPLY QUAD
fSMULD	11 0100	0 0110 1001	MULTIPLY SINGLE TO DOUBLE
fDMULq	11 0100	0 0110 1110	MULTIPLY DOUBLE TO QUAD
FDIVs	11 0100	0 0100 1101	DIVIDE SINGLE
FDIVd	11 0100	0 0100 1110	DIVIDE DOUBLE
FDIVq	11 0100	0 0100 1111	DIVIDE QUAD

FIG. 3A (PRIOR ART)

	Docum ent ID	U	Title	Current OR
1	US 20040 06871 7 A1	<input type="checkbox"/>	System and method of utilizing a hardware component to execute an interpretive language	717/142
2	US 20040 06224 5 A1	<input checked="" type="checkbox"/>	TCP/IP offload device	370/392
3	US 20040 05452 2 A1	<input checked="" type="checkbox"/>	System and method to access web resources from wireless devices	704/8
4	US 20040 04977 4 A1	<input checked="" type="checkbox"/>	REMOTE DIRECT MEMORY ACCESS ENABLED NETWORK INTERFACE CONTROLLER SWITCHOVER AND SWITCHBACK SUPPORT	719/312
5	US 20040 04960 0 A1	<input checked="" type="checkbox"/>	Memory management offload for RDMA enabled network adapters	709/250
6	US 20040 04736 6 A1	<input checked="" type="checkbox"/>	Method for dynamic flow mapping in a wireless network	370/466
7	US 20040 03479 7 A1	<input checked="" type="checkbox"/>	Domain-less service selection	713/201
8	US 20040 01989 1 A1	<input checked="" type="checkbox"/>	Method and apparatus for optimizing performance in a multi-processing system	718/102
9	US 20040 01588 8 A1	<input checked="" type="checkbox"/>	Processor system including dynamic translation facility, binary translation program that runs in computer having processor system implemented therein, and semiconductor device having processor system implemented therein	717/136
10	US 20040 01067 9 A1	<input checked="" type="checkbox"/>	Reducing processor energy consumption by controlling processor resources	713/1
11	US 20030 21717 4 A1	<input checked="" type="checkbox"/>	Establishing an IP session between a host using SIP and a device without an IP address	709/237
12	US 20030 21266 0 A1	<input checked="" type="checkbox"/>	Database scattering system	707/1
13	US 20030 20448 2 A1	<input checked="" type="checkbox"/>	Data search system	707/1
14	US 20030 19822 6 A1	<input checked="" type="checkbox"/>	SEGMENTATION PROTOCOL THAT SUPPORTS COMPRESSED SEGMENTATION HEADERS	370/393
15	US 20030 15430 7 A1	<input checked="" type="checkbox"/>	Method and apparatus for aggregate network address routes	709/245
16	US 20030 11279 4 A1	<input checked="" type="checkbox"/>	System and method for multiple PDP contexts with a single PDP address at a GGSN	370/352
17	US 20030 10597 6 A1	<input checked="" type="checkbox"/>	Flow-based detection of network intrusions	713/201

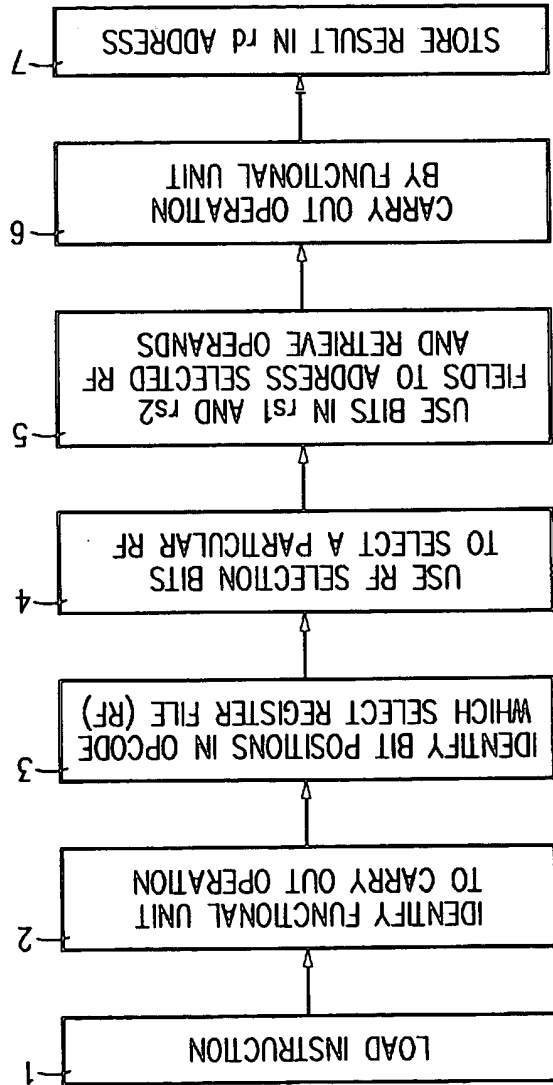


FIG. 7

	Docum ent ID	U	Title	Current OR
18	US 20030 09754 8 A1	<input checked="" type="checkbox"/>	Context execution in pipelined computer processor	712/228
19	US 20030 09754 7 A1	<input checked="" type="checkbox"/>	Context scheduling	712/228
20	US 20030 08433 7 A1	<input checked="" type="checkbox"/>	Remotely controlled failsafe boot mechanism and manager for a network device	713/200
21	US 20030 08158 2 A1	<input checked="" type="checkbox"/>	Aggregating multiple wireless communication channels for high data rate transfers	370/338
22	US 20030 07184 7 A1	<input checked="" type="checkbox"/>	Notification of messages to a terminal by means of vector image	345/772
23	US 20030 06021 0 A1	<input checked="" type="checkbox"/>	System and method for providing real-time and non-real-time services over a communications system	455/452 .1
24	US 20030 05775 1 A1	<input checked="" type="checkbox"/>	Dude tandem	297/243
25	US 20030 03354 1 A1	<input checked="" type="checkbox"/>	Method and apparatus for detecting improper intrusions from a network into information systems	713/201
26	US 20030 02623 0 A1	<input checked="" type="checkbox"/>	Proxy duplicate address detection for dynamic address allocation	370/338
27	US 20030 01447 3 A1	<input checked="" type="checkbox"/>	Multi-thread executing method and parallel processing system	718/107
28	US 20030 01447 1 A1	<input checked="" type="checkbox"/>	Multi-thread execution method and parallel processor system	718/107
29	US 20030 00514 4 A1	<input checked="" type="checkbox"/>	EFFICIENT CLASSIFICATION MANIPULATION AND CONTROL OF NETWORK TRANSMISSIONS BY ASSOCIATING NETWORK FLOWS WITH RULE BASED FUNCTIONS	709/235
30	US 20020 19161 2 A1	<input checked="" type="checkbox"/>	Method and apparatus for automatically determining an appropriate transmission method in a network	370/392
31	US 20020 18778 8 A1	<input checked="" type="checkbox"/>	Mobile communications system	455/450
32	US 20020 18144 8 A1	<input checked="" type="checkbox"/>	Prevention of spoofing in telecommunications systems	370/352
33	US 20020 13620 6 A1	<input checked="" type="checkbox"/>	Recursive query for communications network data	370/352
34	US 20020 10796 4 A1	<input checked="" type="checkbox"/>	Mobile communication system and data transferring method for use with mobile communication system	709/225

I

BACKGROUND

The following background is in the context of a processor

code) to format for all OMASIMC processors. The op code 10 contains various fields including an operation field (opt) identifying

identifying the address in a register all where the result of the operation is to be stored, and a general class field (op3)

integer operation, or a load/store operation. Bit positions 30 and 31 identify a larger class. The bit positions of the various

variety of architectures. Additional information regarding the UltraSPARC™ op code format and structure is found in the UltraSPARC™ Architecture Manual.

It is very important that the various field lengths, op code organization, and instruction sets for a particular architecture-

The op code 10 identifies that a particular functional unit is to perform the operation in field op1 on the operands

Hence, the register file can only have 32 addressable locations. In the UltraSPARC™ processor, each word in a

multiplier and a VIS functional unit, require the full 32 addressable locations in the register file for worst case accesses.

number of operations simultaneously to increase the overall speed of the processor. However, the optimal performance

SUMMARY

shown in FIG. 1. Applicants have identified certain bit positions in the op

dance with one embodiment of the invention.

DETAILED DESCRIPTION OF THE
PREFERRED EMBODIMENTS

typically for generating a graphics display created by a computer. FIG. 2 shows the alpha-numeric op code, the

storage location for the result of the operation are identified in the rs1, rs2, and rd fields in the op code format of FIG. 1,

previously described. The fact that the op field identifies VIS instructions is conveyed in the op3 field (FIG. 1) by the

	Docum ent ID	U	Title	Current OR
35	US 20020 09984 4 A1	<input checked="" type="checkbox"/>	Load balancing and dynamic control of multiple data streams in a network	709/232
36	US 20020 08784 3 A1	<input checked="" type="checkbox"/>	Method and apparatus for reducing components necessary for instruction pointer generation in a simultaneous multithreaded processor	712/228
37	US 20020 08757 3 A1	<input checked="" type="checkbox"/>	Automated prospector and targeted advertisement assembly and delivery system	707/102
38	US 20020 06243 5 A1	<input checked="" type="checkbox"/>	PRIORITIZED INSTRUCTION SCHEDULING FOR MULTI-STREAMING PROCESSORS	712/7
39	US 20020 04282 1 A1	<input checked="" type="checkbox"/>	System and method for monitoring and analyzing internet traffic	709/223
40	US 20020 01071 0 A1	<input checked="" type="checkbox"/>	Method for characterizing a complex system	715/500
41	US 20010 05645 6 A1	<input checked="" type="checkbox"/>	PRIORITY BASED SIMULTANEOUS MULTI-THREADING	718/103
42	US 20010 04868 0 A1	<input checked="" type="checkbox"/>	Method and apparatus for packet transmission with header compression	370/389
43	US 20010 03744 8 A1	<input checked="" type="checkbox"/>	Input replicator for interrupts in a simultaneous and redundantly threaded processor	712/244
44	US 20010 03744 7 A1	<input checked="" type="checkbox"/>	Simultaneous and redundantly threaded processor branch outcome queue	712/239
45	US 20010 03625 5 A1	<input checked="" type="checkbox"/>	Methods and apparatus for providing speech recognition services to communication system users	379/88. 01
46	US 20010 02947 8 A1	<input checked="" type="checkbox"/>	System and method for supporting online auctions	705/37
47	US 20010 01005 3 A1	<input checked="" type="checkbox"/>	Service framework for a distributed object network system	718/105
48	US 67218 06 B2	<input checked="" type="checkbox"/>	Remote direct memory access enabled network interface controller switchover and switchback support	719/312
49	US 67213 33 B1	<input checked="" type="checkbox"/>	Point to point protocol multiplexing/demultiplexing method and apparatus	370/469
50	US 67078 13 B1	<input checked="" type="checkbox"/>	Method of call control to minimize delays in launching multimedia or voice calls in a packet-switched radio telecommunications network	370/356
51	US 66584 47 B2	<input checked="" type="checkbox"/>	Priority based simultaneous multi-threading	718/103
52	US 66256 35 B1	<input checked="" type="checkbox"/>	Deterministic and preemptive thread scheduling and its use in debugging multithreaded applications	718/102
53	US 65642 67 B1	<input checked="" type="checkbox"/>	Network adapter with large frame transfer emulation	709/250

FIG. 3A illustrates a partial instruction set for a floating-point multiply and divide functional unit which identifies the alpha-numeric op code, the code in the op3 field for identifying the functional unit, the code in the op1 field for identifying the particular operation to be performed, and a description of the corresponding operation.

FIG. 3B is a more complete instruction set for the floating point operations, identifying the op code nomenclature as well as the corresponding 9-bit op code (8:0) in bit positions 13-10 in FIG. 1.

Applicants have noted that the first two bits (bit positions 12 and 13) at the left of the op code in the VIS of FIG. 2 are always 00 and that the first and fifth bits (bit positions 9 and 13) from the left of the op code in the floating point instruction set are always 00. Hence, once it is identified that the instruction set is either VIS (op3=36₁₆) or a floating point (op3=34₁₆) operation, it is then known that the bits at bit positions 12 and 13 for a VIS operation and bit positions 9 and 13 for a floating point operation convey no information. Though a similar situation applies with other instruction sets either in the operation field or another field, only a VIS addition/subtraction instruction set and a floating-point instruction set are presented for illustration.

FIG. 4 illustrates an op code format 11 similar to that of FIG. 1 but modified to identify the two bit positions 12 and 13 in the operation field op1 for a VIS operation, which now contains bits for identifying a particular register file among a plurality of register files.

FIG. 5 illustrates a similar op code format 16 for a floating point operation, which identifies the two bit positions 9 and 13 in the op1 field used to identify a particular register file.

Since two bits in the VIS and floating point op code are available for designating a register file, due to these two bits being always the same, one of up to four register files may be identified in these two bit positions. Thus, since up to four register files, each having 32 addresses, can be identified by the two bits and each of the rs1, rs2, and rd fields may identify one out of 32 addresses, the registers have been increased from 32 to up to 128.

An application of this technique will now be described with respect to FIGS. 6 and 7. FIG. 6 illustrates modifications to a conventional processor for implementing the present invention. Since the architecture of processors is well known to those skilled in the art, a detailed description of the operation of processors is unnecessary to fully explain the present invention.

In FIG. 6, data is stored in designated addresses in a group of register files 18, 19, 20 and 21 via a data cache 22. An instruction in the form of the op code format 11 or 16 in FIGS. 4 or 5 is provided from an external RAM 24 into an instruction-cache 30.

The instruction in the instruction cache 30 is then provided to an instruction queue 32 which in one embodiment has a capacity for 16 instructions. The next instruction in the queue 32, under the control of an instruction queue controller 34, is applied to an FP/VIS instruction controller 36, as identified in step 1 of FIG. 7. Controller 36 then dispatches the instruction to the proper functional unit 40-47 (step 2 in FIG. 7). Examples of functional units are illustrated as floating point and VIS adders 40-43 and floating point and VIS multipliers 44-47. Additional functional units would typically exist. Controller 36 determines the availability of the various functional units prior to tasking a particular functional unit to carry out the instruction. Controller 36 is typically a hardwired logic circuit.

Controller 36 addresses the register files to retrieve the operands for the designated functional unit as follows.

FIG. 3A illustrates a partial instruction set for a floating-point multiply and divide functional unit which identifies the alpha-numeric op code, the code in the op3 field for identifying the functional unit, the code in the op1 field for identifying the particular operation to be performed, and a description of the corresponding operation.

FIG. 3B is a more complete instruction set for the floating point operations, identifying the op code nomenclature as well as the corresponding 9-bit op code (8:0) in bit positions 13-10 in FIG. 1.

Applicants have noted that the first two bits (bit positions 12 and 13) at the left of the op code in the VIS of FIG. 2 are always 00 and that the first and fifth bits (bit positions 9 and 13) from the left of the op code in the floating point instruction set are always 00. Hence, once it is identified that the instruction set is either VIS (op3=36₁₆) or a floating point (op3=34₁₆) operation, it is then known that the bits at bit positions 12 and 13 for a VIS operation and bit positions 9 and 13 for a floating point operation convey no information. Though a similar situation applies with other instruction sets either in the operation field or another field, only a VIS addition/subtraction instruction set and a floating-point instruction set are presented for illustration.

FIG. 4 illustrates an op code format 11 similar to that of FIG. 1 but modified to identify the two bit positions 12 and 13 in the operation field op1 for a VIS operation, which now contains bits for identifying a particular register file among a plurality of register files.

FIG. 5 illustrates a similar op code format 16 for a floating point operation, which identifies the two bit positions 9 and 13 in the op1 field used to identify a particular register file.

Since two bits in the VIS and floating point op code are available for designating a register file, due to these two bits being always the same, one of up to four register files may be identified in these two bit positions. Thus, since up to four register files, each having 32 addresses, can be identified by the two bits and each of the rs1, rs2, and rd fields may identify one out of 32 addresses, the registers have been increased from 32 to up to 128.

An application of this technique will now be described with respect to FIGS. 6 and 7. FIG. 6 illustrates modifications to a conventional processor for implementing the present invention. Since the architecture of processors is well known to those skilled in the art, a detailed description of the operation of processors is unnecessary to fully explain the present invention.

Controller 36 addresses the register files to retrieve the operands for the designated functional unit as follows.

	Docum ent ID	U	Title	Current OR
54	US 65325 54 B1	<input checked="" type="checkbox"/>	Network event correlation system using formally specified models of protocol behavior	714/43
55	US 65196 36 B2	<input checked="" type="checkbox"/>	Efficient classification, manipulation, and control of network transmissions by associating network flows with rule based functions	709/223
56	US H0020 51 H	<input checked="" type="checkbox"/>	System and method for providing multiple quality of service classes	370/395 .21
57	US 64775 62 B2	<input checked="" type="checkbox"/>	Prioritized instruction scheduling for multi-streaming processors	718/108
58	US 64704 43 B1	<input checked="" type="checkbox"/>	Pipelined multi-thread processor selecting thread instruction in inter-stage buffer based on count information	712/205
59	US 64346 76 B1	<input checked="" type="checkbox"/>	FIFO with random re-read support and its application	711/154
60	US 64337 95 B1	<input checked="" type="checkbox"/>	System for integrating an on-line service community with a foreign service	345/738
61	US 63634 24 B1	<input checked="" type="checkbox"/>	Reuse of services between different domains using state machine mapping techniques	709/224
62	US 63380 78 B1	<input checked="" type="checkbox"/>	System and method for sequencing packets for multiprocessor parallelization in a computer network system	718/102
63	US 62956 00 B1	<input checked="" type="checkbox"/>	Thread switch on blocked load or store using instruction thread field	712/228
64	US 62955 57 B1	<input checked="" type="checkbox"/>	Apparatus for simulating internet traffic	709/224
65	US 62928 88 B1	<input checked="" type="checkbox"/>	Register transfer unit for electronic processor	712/225
66	US 62721 48 B1	<input checked="" type="checkbox"/>	Scheme for reliable communications via radio and wire networks using transport layer connection	370/469
67	US 62333 15 B1	<input checked="" type="checkbox"/>	Methods and apparatus for increasing the utility and interoperability of peripheral devices in communications systems	379/88. 01
68	US 62298 80 B1	<input checked="" type="checkbox"/>	Methods and apparatus for efficiently providing a communication system with speech recognition capabilities	379/88. 01
69	US 62232 74 B1	<input checked="" type="checkbox"/>	Power-and speed-efficient data storage/transfer architecture models and design methodologies for programmable or reusable multi-media processors	712/34
70	US 62162 20 B1	<input checked="" type="checkbox"/>	Multithreaded data processing method with long latency subinstructions	712/219
71	US 62090 18 B1	<input checked="" type="checkbox"/>	Service framework for a distributed object network system	718/105
72	US 61547 77 A	<input checked="" type="checkbox"/>	System for context-dependent name resolution	709/227
73	US 61051 27 A	<input checked="" type="checkbox"/>	Multithreaded processor for processing multiple instruction streams independently of each other by flexibly controlling throughput in each instruction stream	712/215
74	US 60731 59 A	<input checked="" type="checkbox"/>	Thread properties attribute vector based thread selection in multithreading processor	718/103
75	US 60527 08 A	<input checked="" type="checkbox"/>	Performance monitoring of thread switch events in a multithreaded processor	718/108
76	US 60208 84 A	<input checked="" type="checkbox"/>	System integrating an on-line service community with a foreign service	345/747

translates into using one-third as many bits as would be required had each of the address fields been expanded by two bits.

Those skilled in the art will understand that the concepts described herein may be applied to any other architecture, such as 16-bit or 64-bit architecture, and any op code format.

While particular embodiments of the present invention have been shown and described, it will be obvious to those skilled in the art that changes and modifications may be made without departing from this invention in its broader aspects and, therefore, the appended claims are to encompass within their scope all such changes and modifications as fall within the true spirit and scope of this invention.

What is claimed is:

1. A method carried out by a processor, said processor carrying out designated operations on at least one operand, after retrieving said at least one operand from one or more registers, and writing a result in a destination register, addresses of said at least one operand in said one or more registers and said destination register being identified by at least two N-bit fields in an operation code format, said method comprising:

identifying in said operation code format a first address field, having N bits, containing a first address of a first storage location within a register file, said register file having up to 2^N addressable locations;

identifying in said operation code format a second address field, having N bits, containing a second address of a second storage location within said register file;

applying said first address and said second address to said selected register file to read said at least one operand from said selected register file and write said result into said selected register file.

2. The method of claim 1 wherein one or more bit positions 12 and 13 in an operation code format for a processor carrying out a visual instruction set operation.

3. The method of claim 1 wherein one or more bit positions 12 in an operation code format for a processor carrying out a visual instruction set operation.

4. The method of claim 1 wherein one or more bit positions 13 in an operation code format for a processor carrying out a visual instruction set operation.

5. The method of claim 1 wherein one or more bit positions 9 and 13 in an operation code format for a processor carrying out a floating point operation.

6. The method of claim 1 wherein one or more bit positions 9 in an operation code format for a processor carrying out a floating point operation.

7. The method of claim 1 wherein one or more bit positions storing said register file selection bits comprise bit

position 13 in an operation code format for a processor carrying out a floating point operation.

8. A method for programming a processor, said processor carrying out designated operations on at least one operand, after retrieving said at least one operand from one or more registers, and writing a result in a destination register, addresses of said at least one operand in said one or more registers and said destination register being identified by at least two N-bit fields in an operation code format, said method comprising:

providing in said operation code format a first address field, having N bits, containing a first address of a first storage location within a register file, said register file having up to 2^N addressable locations;

providing in said operation code format a second address field, having N bits, containing a second address of a second storage location within said register file;

providing in said operation code format one or more registers, and writing a result in a destination register, addresses of said at least one operand in said one or more registers and said destination register being identified by at least two N-bit fields in an operation code format, said method comprising:

providing an instruction in said operation code format for causing said processor to apply said first address and said second address to said selected register file and write said result into said selected register file.

9. The method of claim 8 wherein one or more bit positions 12 and 13 in an operation code format for a processor carrying out a visual instruction set operation.

10. The method of claim 8 wherein one or more bit positions 12 in an operation code format for a processor carrying out a visual instruction set operation.

11. The method of claim 8 wherein one or more bit positions 13 in an operation code format for a processor carrying out a floating point operation.

12. The method of claim 8 wherein one or more bit positions storing said register file selection bits comprise bit positions 9 and 13 in an operation code format for a processor carrying out a floating point operation.

13. The method of claim 8 wherein one or more bit positions storing said register file selection bits comprise bit position 9 in an operation code format for a processor carrying out a floating point operation.

14. The method of claim 8 wherein one or more bit positions storing said register file selection bits comprise bit position 13 in an operation code format for a processor carrying out a floating point operation.

15. A processor comprising:

a plurality of functional units, each functional unit for carrying out a set of instructions on one or more operands and storing a result of an operation in a register;

a plurality of register files, each register file having up to 2^N addressable locations;

an input port receiving an instruction in an operation code format, said instruction including at least one operand address field containing a first address having N bits designating an address of a first storage location within

	Docum ent ID	U	Title	Current OR
77	US 59000 25 A	<input checked="" type="checkbox"/>	Processor having a hierarchical control register file and methods for operating the same	712/248
78	US 58812 77 A	<input checked="" type="checkbox"/>	Pipelined microprocessor with branch misprediction cache circuits, systems and methods	712/239
79	US 58549 13 A	<input checked="" type="checkbox"/>	Microprocessor with an architecture mode control capable of supporting extensions of two distinct instruction-set architectures	712/210
80	US 58128 11 A	<input checked="" type="checkbox"/>	Executing speculative parallel instructions threads with forking and inter-thread communication	712/216
81	US 57963 93 A	<input checked="" type="checkbox"/>	System for intergrating an on-line service community with a foreign service	345/733
82	US 57245 65 A	<input checked="" type="checkbox"/>	Method and system for processing first and second sets of instructions by first and second types of processing systems	712/245
83	US 53576 17 A	<input type="checkbox"/>	Method and apparatus for substantially concurrent multiple instruction thread processing by a single pipeline processor	712/245

a register file, said instruction also including a designation address field containing a second address having N bits designating a second storage location for a result of an operation carried out by a designated functional unit, said instruction also containing one or more register file selection bits outside of said at least one operand address field and said destination address field for designating one of a plurality of register files as a selected register file for said first address and said second address, each register file in said plurality of register files having up to 2^N addressable locations; and a controller receiving said instruction and selecting one of said plurality of register files as a selected register file

based on said one or more register file selection bits, said controller applying said first address to said selected register file to read an operand from said selected register file to apply said operand to a designated functional unit.
 16. The processor of claim 15 wherein said designated functional unit writes said result into said second storage location in said selected register file after performing an operation on one or more operands from said selected register file.

* * * * *